

ShowCase®
a division of SPSS

STRATEGY™ 4.5
Administrator's Guide

18794 001

For more information about SPSS[®] software products, please visit our Web site at <http://www.spss.com> or contact

SPSS Inc.
233 South Wacker Drive, 11th Floor
Chicago, IL 60606-6412
Tel: (312) 651-3000
Fax: (312) 651-3668

SPSS is a registered trademark and its other product names are the trademarks of SPSS Inc. for its proprietary computer software. No material describing such software may be produced or distributed without the written permission of the owners of the trademark and license rights in the software and the copyrights in the published materials.

The SOFTWARE and documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subdivision (c)(1)(ii) of The Rights in Technical Data and Computer Software clause at 52.227-7013. Contractor/manufacturer is SPSS Inc., 233 South Wacker Drive, 11th Floor, Chicago, IL 60606-6412.

Acrobat is a registered trademark of Adobe Systems Inc.
AS/400 is a registered trademark of IBM Corporation.
Essbase is a registered trademark of Hyperion Solutions Corporation.
Orion is a trademark of Evermind Data HB.
ShowCase is a registered trademark of ShowCase Corporation, a division of SPSS Inc.
Windows is a registered trademark of Microsoft Corporation.

General notice: Other product names mentioned herein are used for identification purposes only and may be trademarks or registered trademarks of their respective companies in the United States and other countries.

STRATEGY[™] 4.5 Administrator's Guide
Copyright © 2001 by SPSS Inc.
All rights reserved.
Printed in the United States of America.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

1 2 3 4 5 6 7 8 9 0 05 04 03 02

ISBN 1-56827-278-2

Preface

Welcome to the *STRATEGY™ 4.5 Administrator's Guide*. This guide discusses the administration and maintenance of the STRATEGY software. For information about installing STRATEGY, see the *STRATEGY 4.5 Installation Guide*.

This preface provides the following information:

- What you'll find in this guide
- Where to find additional information
- How to find documentation online
- How to contact the documentation team to provide feedback about this guide

Adobe Acrobat (.pdf) versions of the STRATEGY guides are available on the STRATEGY CD-ROM. Viewing and printing the documents requires Adobe Acrobat Reader. If necessary, you can download it at no cost from www.adobe.com, or you can install it from the CD-ROM. Click Browse the CD-ROM, and open the Tools folder. Use the Adobe Acrobat Reader online Help to get answers to your questions about viewing and navigating the documents.

About This Guide

- **Chapter 1, "Technical Support,"** contains ShowCase contact information.
- **Chapter 2, "Starting and Stopping the Servers,"** explains how to start and stop the Warehouse Manager and Essbase/400 servers and how to prepare the servers for client connections.
- **Chapter 3, "Controlling and Managing Performance,"** explains how to manage performance to ensure successful deployment.
- **Chapter 4, "Managing STRATEGY Data on the iSeries™ (AS/400®),"** describes how to maintain your Warehouse Manager library, including backing up and restoring files, checking for usage, and configuring TCP/IP.
- **Chapter 5, "Warehouse Manager Server Maintenance,"** describes typical server maintenance procedures.
- **Chapter 6, "Essbase/400 Server Maintenance,"** explains how to use and maintain your Essbase/400 library. If you are familiar with Essbase on other platforms, we recommend reading this chapter to understand how Essbase on an iSeries platform differs.

- **Appendix A, "Commands,"** contains a list of commands commonly used to maintain iSeries libraries.
- **Appendix B, "Sample Databases,"** describes the sample databases installed with Warehouse Manager Server.

Additional Resources

- For getting started and deployment information, see the *STRATEGY 4.5 Installation Guide*.
- For information about how STRATEGY licensing and passwords work, see the *STRATEGY 4.5 Installation Guide*.
- For information about installing STRATEGY PC software, see the *STRATEGY 4.5 Installation Guide*.
- For tasks and steps for product features unrelated to the maintenance of STRATEGY iSeries software (Warehouse Manager Server and Essbase/400), see the online Help in each application or the *Quick Start* guides on the Support Web site at <http://support.spss.com>. After you log on, click Documentation and then STRATEGY Manuals.
- For information about configuring TCP/IP, see the *STRATEGY 4.5 Installation Guide*.
- For information about upgrading to STRATEGY 4.5, see the *STRATEGY 4.5 Upgrade Guide* at <http://support.spss.com/upgradeinfo/>.

Finding Documentation Online

You can view and download additional documentation after you install STRATEGY. You must have an Internet connection to view the documents. From the Windows Start menu, choose:

Programs
ShowCase STRATEGY
Documentation

Send Us Your Comments

Your comments are important to us. Please tell us what you think about the STRATEGY documentation by contacting:

Documentation Team
ShowCase®, a division of SPSS Inc.
4115 Highway 52 NW, Suite 300
Rochester, MN 55901-0144
Tel: (507) 288-5922
E-mail: docfeedback@spss.com

Customer Service

If you have any questions concerning your shipment or account, contact your local office, listed on p. v. Please have your serial number ready for identification when calling.

ShowCase
A Division of SPSS
233 S. Wacker Drive.
11th Floor
Chicago, IL 60606
TEL +312.651.3000

ShowCase UK Limited
4 River Court, Albert Drive
Woking, Surrey GU21 5RP
TEL +44 (0) 1483 744300

ShowCase France
22-24 Avenue Hoche
75008 Paris
TEL +33 (0) 1 56 69 18 18

ShowCase Benelux
Buro & Design Center, B38
Heizel Esplanade Heysel
B-1020 Brussels
TEL +32 2 474 03 60

ShowCase Deutschland GmbH
Mergenthalerallee 79-81
D-65760 Eschborn
TEL +49 6196 6510 0

ShowCase Nederland
Lichtenauerlaan 102-120
Postbus 4365
3006 AJ Rotterdam
TEL +31 (0) 10 204 57 40

Spain SPSS
Luchana 23
5 Planta
28010, Madrid
TEL +39 91 443 37 00

SPSS Japan Inc
ShowCase Division
10/F Ebisu Prime Square,
1-1-39 Hiroo, Shibuya-Ku,
Tokyo 150
TEL +81 3-5466-5512

SPSS Australasia Pty Ltd
ShowCase Division
Level 9, 181 Miller Street
North Sydney, Australia
NSW 2060
TEL +61 2 9954 5560

Sweden AB
(formerly SPSS Scandinavia)
Target 1
Sundbyberg, Sweden
SE-17267
TEL +46 8 506 105 50

SPSS Singapore Pte. Ltd.
21 Cantonment Road
The Courtyard, Level 1
Singapore, Singapore
089742
TEL +65 324 5150

Contents

1	<i>Technical Support</i>	1
	Contacting Support	1
	Internet Support.	1
	North American Support.	2
	European Support.	3
	Japan and Asia Pacific Support	3
2	<i>Starting and Stopping the Servers</i>	5
	Warehouse Manager Server (Including Analyzer OLAP Functionality)	5
	Starting Warehouse Manager Server	5
	Stopping Warehouse Manager Server	6
	Essbase/400 Server.	6
	Starting Essbase/400 Server	6
	Stopping Essbase/400 Server	7
	Orion Server	7
3	<i>Controlling and Managing Performance</i>	9
	Managing a Relational Database	9
	Planning Security.	9
	Controlling Query/Report Writer System Usage	18
	Controlling Level of Parallelism.	21
	Enabling Expert Cache.	22
	Tuning System Performance	22
	DB2 Symmetric Multiprocessing for OS/400	23
	Using Auditing to Understand Your Queries.	23
	Improving Query Performance	28
	Optimizing Performance for Warehouse Builder	32
	Managing a Multidimensional Database	36
	Optimizing Essbase/400 Performance	36

4 Managing STRATEGY Data on the iSeries (AS/400) 41

Moving Data from a Test to a Production Environment	41
Saving and Restoring Data	41
Redirecting Data	42
Creating Custom Library Lists for Users	43
Backing Up Data	43
Temporary Backups	44
Disaster Recovery Backups	44
After Restoring Server Information	46
Changing Physical Files That Have Dependent Data Views	47
Changing Data Views That Have Dependent Data Views	47
Managing Passwords	48
Changing Passwords with a STRATEGY Product	48
Updating Passwords with the SCCHKPWD Command	48
Managing Information Workflow	50
Running Distribution Sets with STRDD and SBMDD	50
Pulling iSeries Data into Essbase NT	50
STRATEGY Integration with J. D. Edwards (JDE)	52
Non-iSeries Profile Sign-On	52
Editing ShowCase.cfg to Use an Alternate F0094 File	53
User-Defined Codes (UDC)	54
Multiple Library List Environments	54
Selecting Group Library Lists	54

5 Warehouse Manager Server Maintenance 57

Using TCP/IP with STRATEGY	57
Configuring Warehouse Manager TCP/IP with STRSCSVR.	57
Finding Available Port Numbers	58
Migrating Users to a New Warehouse Manager Library.	58
Checking Usage of a Warehouse Manager Library	58
Listing Connections to Relational Applications Only	58
Listing All Connections to All STRATEGY Applications	59
Ending STRATEGY TCP/IP (Stopping the Server).	60
Uninstalling a Warehouse Manager Library	60
Uninstalling Warehouse Manager Server Libraries.	60

Uninstalling Warehouse Manager Server Libraries (pre-STRATEGY)	61
Backing Up Warehouse Manager Server Information	61
Saving STRATEGY iSeries Media and Passwords	61
Saving Server Information.	61
Restoring Warehouse Manager Server Information	63
Installing STRATEGY iSeries Media and Passwords.	63
Restoring Server Information	63
Detecting and Removing Obsolete Information.	65
Restricting Access to a Warehouse Manager Library	65
Applying the Restricted State	65
Removing the Restricted State	67
Working with Journals	67
Warehouse Manager and iSeries Work Management.	68
Work Management Controls	68
Security Issues	71
Command Authorities	71
User Exit Programs	73
Distributed Relational Database Architecture (DRDA).	74
Analyzer OLAP Maintenance	75
Configuring the iSeries HTTP Server.	75
Using National Language Support	82
HTML Documents.	83
Getting Users Started	84
About Distributed Archive Files.	88

6 *Essbase/400 Server Maintenance* **91**

Understanding the Essbase/400 Configuration	91
Essbase/400 Installation Library and Path.	91
Essbase/400 QSYS.LIB Objects Created During Installation	91
Using Essbase/400 Configuration File Settings	92
Using Configuration Files from a 5250 Display Session	93
Using Essbase/400 Environment Settings	93
Understanding the Configuration and Environment Relationship.	95
Configuring Essbase/400 for Java	95
Using Essbase/400 National Language Support	95
Migrating to the Euro CCSID	96
Displaying the Current EBCDIC CCSID	96
Displaying the Current ASCII CCSID	97
Changing the ASCII and EBCDIC CCSIDs	97

Administering Essbase/400	98
Command Interfaces	98
Using an Input Script File	98
Using the Command Line Interface	101
ESSCMD/400 File Differences	101
Getting Help with ESSCMD or MaxL Commands	103
Using the ESSCMD API Toolkit.	103
Overview	103
Command Structure	104
Auto-logout Considerations.	105
Supplemental Commands.	105
Output Processing	107
Controlling Error Handling When Programming	107
Debugging	108
Maintaining Essbase/400 Log Files	108
Index Cache Size.	109
Loading DB2/400 Data into Essbase/400	109
Setting Up Your iSeries to Allow SQL Loads	109
Loading Data Using the SQL Interface Support	110
Substitution Variables in Rule File SQL Statements	110
Flat File Support	111
Using Warehouse Builder with Essbase/400	111
Dimension Build and Data Load Rejected Records	112
Reloading Rejected Records	112
Naming Essbase/400 Server	112
Security	112
ALLOBJ Authority	113
Integrated Security	113
Setting Up Essbase/400 Integrated Security	114
Setting Up Group-Level Control	119
Essbase/400 Jobs	123
SQL Drill-Through Using Essbase Add-ins	124
Message Logging	124
Backing Up Essbase/400 Files	124
Restoring Essbase/400 Files	125

Appendixes

A Commands **127**

iSeries Commands	127
Warehouse Manager Server Commands	128
CHGSCSVRA Analyzer OLAP-Specific Settings	130
SCCHKPWD Command Parameters	134
Essbase/400 Commands	136

B Sample Databases **139**

Installing the Sample Databases	139
SCSAMPLE45	139
SCDB45	139
Contents of the Sample Databases	140
SCSAMPLE45 Database	141
SCDB45 Database	150
Installing the Sample Library after Installation with RSTSMPLB	162

Index **163**

Product Enhancement Request Form **171**

Technical Support

Contacting Support

Phone. Our goal is to provide one-on-one technical support. However, based on the complexity and/or frequency of calls, you may reach our phone-mail system. Leave your name, company name, and telephone number and a support member will call you back.

Fax. Send your name, company name, and telephone and fax numbers along with the details of your request and the name of the product with which you need assistance.

Internet. Enter your questions and requests on the Support Web site. You can also use the site to find documentation, download patches, open a case, and participate in the customer forum. See “Internet Support” below.

E-mail. Include your name, e-mail address, company name, the feature that you have a question about, and your question or request.

Internet Support

Table 1-1
Internet addresses

Site	Address
Support Web site	http://support.spss.com
ShowCase, a division of SPSS, Web site	http://www.spssshowcase.com
Documentation team e-mail	docfeedback@spss.com

The following information is on the SPSS Support Web site at <http://support.spss.com>. To access STRATEGY support, click the ShowCase Division Products link.

Frequently asked questions. Find answers to some of the most common questions.

Frequently used resolutions. View the top resolutions of the week, month, or all time periods.

STRATEGY manuals. View and download the latest documentation, including the *Installation Guide*, *Administrator's Guide*, *Upgrade Guide*, and *Quick Start* guides.

Technical bulletins and newsletters. View and download technical bulletins or browse the newsletter archive for technical issues.

Patches. Download patches for all supported STRATEGY products.

Case information. Open a support case and view the status of your cases.

STRATEGY customer forum. Network with other maintenance customers in this discussion group.

Contact information. Find current phone and fax numbers and e-mail addresses for North America, Europe, and Japan/Asia Pacific.

North American Support

Table 1-2
North American services and support information

Services and Support	Phone	Fax	E-mail
North American support (relational products)	(507) 288-5922 Option 2, then option 1	(507) 287-2809	scsupport@spss.com
North American support (multidimensional products)	(507) 288-5922 Option 2, then option 2	(507) 287-2809	scsupport@spss.com
Product training and consulting services	(312) 261-6544	(847) 779-6544	TrainingServices@spss.com

European Support

Table 1-3
European support information

Country	Language	Phone	Fax	E-mail
Austria	German	49 89 48 9074 333	49 89 44 83115	suppge@spss.de
Belgium	Dutch	32 2 477 17 00	32 2 474 03 63	eursupport@spss.com
Belgium	French	32 2 477 17 01	32 2 474 03 63	eursupport@spss.com
Belgium	German	32 2 477 17 02	32 2 474 03 63	eursupport@spss.com
Denmark	English	45 361 58731	32 2 474 03 63	eursupport@spss.com
France	French	33 155 21 25 32	32 2 474 03 63	eursupport@spss.com
Germany	German	49 89 48 9074 333	49 89 44 83115	suppge@spss.de
Italy	Italian and English	39 0280 663036	32 2 474 03 63	eursupport@spss.com
Luxembourg	French	32 2 477 17 03	32 2 474 03 63	eursupport@spss.com
Switzerland	French	41 227 100034	32 2 474 03 63	eursupport@spss.com
Switzerland	German	49 89 48 9074 333	49 89 44 83115	suppge@spss.de
The Netherlands	Dutch	31 205 82 94 96	32 2 474 03 63	eursupport@spss.com
United Kingdom	English	44 1483 719 202	44 1483 719 290	support@spss.co.uk

Japan and Asia Pacific Support

Table 1-4
Japan and Asia Pacific support information

Region	Phone	Fax	E-mail
Japan and Asia Pacific	81 3 5466 7187	81 3 5466 5622	jpscsupport@spss.com

Starting and Stopping the Servers

This chapter outlines the procedures for starting and stopping Warehouse Manager Server and Essbase/400 Server.

Note: Before performing a task on the iSeries (AS/400) that may jeopardize your system stability, such as applying PTFs, it is recommended that you stop the servers.

Warehouse Manager Server (Including Analyzer OLAP Functionality)

This section outlines the commands to start and stop Warehouse Manager Server.

Note: Users of STRATEGY 3.5 and earlier releases will notice that since release 4.1, AOS no longer has its own “start” and “stop” server commands. To simplify managing the servers, the Analyzer OLAP functionality has been integrated into one Warehouse Manager Server. Now, when you start and stop Warehouse Manager Server, you are also starting and stopping Analyzer OLAP features.

Starting Warehouse Manager Server

Running STRSCSVR will start Warehouse Manager Server and enable Analyzer OLAP features that you are licensed to use. If you are using the TCP/IP communications protocol, the Warehouse Manager TCP/IP Server job should start automatically when the installation is finished, provided that TCP/IP was configured correctly.

However, if you opted to set up Warehouse Manager TCP/IP after installation or if you stopped Warehouse Manager Server, you must start Warehouse Manager TCP/IP Server.

- ▶ Log on with QSECOFR or equivalent authority (*JOBCTL authority is required).
- ▶ Add the Warehouse Manager Server library to your library list. This is the name of the library in which the server is installed.
- ▶ Run the Warehouse Manager command serverlib/STRSCSVR, where serverlib is the name of the library in which the server is installed.

Stopping Warehouse Manager Server

To stop the Warehouse Manager Server and Warehouse Manager TCP/IP Server connections:

- ▶ Log on with QSECOFR or equivalent authority (*ALLOBJ and *JOBCTL authorities are required).
- ▶ Add the Warehouse Manager Server library to your library list.
- ▶ Run the Warehouse Manager command `serverlib/ENDSCSVR`, where `serverlib` is the name of the library in which the servers are installed.

The `ENDSCSVR` command will terminate all current TCP/IP connections to the specified library, including Analyzer OLAP functionality. For more information on this and related topics, see “Uninstalling a Warehouse Manager Library” on p. 60 in Chapter 5.

Essbase/400 Server

This section outlines the commands to start and stop Essbase/400 Server.

Starting Essbase/400 Server

Note: Be sure that iSeries TCP/IP (STRTCP) is running *before* starting Essbase/400 Server.

To start Essbase/400, enter the command `essbaselib/STRESSSVR`, where `essbaselib` is the name of the library in which Essbase/400 is installed (the default is `ESSBASE`). This command has one parameter—the job description under which Essbase/400 will run (the default is `*USRPRF`).

If this is the first time that Essbase/400 is being run and these parameters were not set up during installation, you will be prompted for the following information:

- ▶ **Enter your company name.** Enter the name of the company to which Essbase/400 is licensed.
- ▶ **Enter your initial system supervisor user ID.** Enter an Essbase/400 system user ID. This user ID is independent of your OS/400 user profiles. It does not need to match an existing OS/400 user profile name, although you may choose a name that does correspond to an iSeries user ID.
- ▶ **Enter the system password.** Enter a password to correspond with the user ID that you specified in the previous step.
- ▶ **Confirm the system password.** Reenter the password and press Enter.

Your Essbase/400 Server should now be running.

Note: The subsystem for Essbase/400 Server is controlled with the job description parameter included in the `STRESSSVR` command. Change the job description to one that is associated with the subsystem that you want to use.

Stopping Essbase/400 Server

To stop the Essbase/400 Server and Essbase/400 TCP/IP Server connections:

- ▶ Log on with QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).

Note: The GRTESSAUT command allows a user who lacks the proper authority to stop or start Essbase/400 Server.

- ▶ Add the Essbase/400 Server library to your library list.
- ▶ Run the Essbase command `essbaselib/ENDESSSVR <username> <password>`, where `essbaselib` is the name of the library in which the server is installed.

The ENDESSSVR command will terminate all current TCP/IP connections to the specified library.

Orion Server

For information about installing, starting, and stopping the Orion Server, see Table A-3 in Appendix A and the *Enterprise Reporting Administrator's Guide*. To access the guide, from the Windows Start menu, choose:

Programs
 ShowCase STRATEGY
 Documentation

and then click the ER Administrator's Guide link.

Controlling and Managing Performance

Managing the performance of relational and multidimensional databases can be a challenge for the most experienced database administrator. This chapter provides insight into performance planning and tuning of both relational and multidimensional databases. It is divided into two sections: managing a relational database and managing a multidimensional database.

Managing a relational database. These sections focus on performance tuning a relational database.

- Planning security
- Controlling Query/Report Writer system usage
- Controlling level of parallelism
- Enabling expert cache
- Tuning system performance
- DB2 symmetric multiprocessing for OS/400
- Using auditing to understand queries
- Improving query performance
- Optimizing performance for Warehouse Builder

Managing a multidimensional database. This section focuses on performance tuning and planning a multidimensional database.

- Optimizing Essbase/400 performance

Managing a Relational Database

This section focuses on tuning your relational database performance.

Planning Security

This section outlines the security issues to consider when deploying STRATEGY for relational data access or multidimensional data analysis.

iSeries (AS/400) Security Considerations

When deploying Warehouse Manager products, consider your current iSeries security environment. Items to address when creating your security plan include:

- What information do you currently store or plan to store on your system?
- Will you use STRATEGY products on an iSeries system that is also currently used for transaction data processing, or will you use STRATEGY software on a new iSeries system that will be dedicated to data warehouse processing?
- Who needs access to the information? Which applications will they use to gain access to the information?
- What is your overall approach to security? Do you have a restrictive or a non-restrictive security approach?
- What type of information will you secure (confidential, competitive, and operations information)?
- What security scheme do you use for client-server applications other than Warehouse Manager products?
- What other vendors' applications affect the way that you set up your system security?
- What iSeries system security level will you use?

It is important to plan your overall iSeries security approach or have a thorough understanding of the existing iSeries security approach. This will allow you to choose the best way to implement Warehouse Manager security to meet your current and future security needs.

Refer to IBM documentation for more detailed iSeries system security information.

Warehouse Manager Server Security

When you set security for Warehouse Manager Server, it will affect all client applications connecting to Warehouse Manager as their server. These include:

- Warehouse Manager Client
- Warehouse Builder
- Query
- Report Writer

Several methods are provided that affect your server security. These methods interact with your current iSeries security settings, your current applications, applications purchased from other vendors, and your additional security measures (such as exit programs) to provide your total security environment. The primary Warehouse Manager security methods include:

- Using Warehouse Manager Client to apply library security, table security, column security, or row security for specific users or groups of users.
- Using Warehouse Manager Client to change server options. These options affect all users.
- Using Warehouse Builder or Query with batch output to a file to create new iSeries objects.

Although the following sections make recommendations for your Warehouse Manager security plan based on specific situations, in general, you will always follow these steps:

- ▶ Designate at least one user profile as a data administrator.
- ▶ Review your user-profile scheme and update it if desired.
- ▶ Decide on your general security policy. Choose whether to restrict global access to files and then authorize specific users to specific files, or to allow global access to files and then exclude specific users from specific files. Implementing this decision requires the following:
 - Set the server options to allow all users to see all libraries and tables, or set the server options to exclude all users from all libraries and tables.
 - Use Warehouse Manager Client to apply authorities to specific tables for specific users.
- ▶ Decide whether to restrict users to libraries within their iSeries library lists.
- ▶ Apply any necessary column filters.
- ▶ Apply any necessary row filters.
- ▶ Ensure the correct users are authorized to, or restricted from, specific STRATEGY applications.
- ▶ If desired, use Warehouse Manager products to output directly to libraries that users are authorized to access.

Designating Data Administrators

Using Warehouse Manager Client, designate one or more user profiles as data administrators. The data administrator(s) will always be able to view all libraries and tables in Warehouse Manager products, and they will also be able to set server options. If you choose a user profile without *ALLOBJ authority to be a data administrator, you must also complete the following steps:

- ▶ Using a profile with *ALLOBJ authority, use Warehouse Manager Client to authorize that user to use Warehouse Manager Client software. See the Warehouse Manager Client online Help for more information.
- ▶ Using a profile with *ALLOBJ authority, give the intended data administrator iSeries *ALL authority to the user profiles that he or she will need to administer. For example, if you want the data administrator to administer all user profiles, use the following iSeries command:

```
GRTOBJAUT OBJ(QSYS/*ALL) OBJTYPE(*USRPRF) USER(userid) AUT(*ALL)
```

- ▶ With a profile that has *ALLOBJ authority, give the intended data administrator iSeries *ALL authority to the libraries and tables that he or she will need to administer. For example, if you want the data administrator to administer all libraries, use the following iSeries command:

```
GRTOBJAUT OBJ(QSYS/*ALL) OBJTYPE(*LIB) USER(userid) AUT(*ALL)
```

- ▶ With a profile that has *ALLOBJ authority, give the intended data administrator iSeries *USE authority to the following commands in the Warehouse Manager Server library on the iSeries:
 - CHGSCAUT
 - CHKSRVRCAT
 - CHGSRVRDFT
 - CHGSRVRSTS
 - DSPSCAUT
 - RMVSCAUT
 - SETLIBACC

iSeries *USE authority allows the data administrator to run these commands and change server options. For example, use the following iSeries command to grant the user authority to CHGSRVRDFT:

```
GRTOBJAUT OBJ(SCSERVER/CHGSRVRDFT) OBJTYPE(*CMD)USER(userid) AUT(*USE)
```

- ▶ With a profile that has *ALLOBJ authority, give the data administrator iSeries *CHANGE authority to the object SCVERSION in the Warehouse Manager Server library on the iSeries. This allows the data administrator to change the security and *ALLOBJ server options. For example, use the following iSeries command:

```
GRTOBJAUT OBJ(SCSERVER/SCVERSION) OBJTYPE(*USRSPC) USER(userid) AUT(*CHANGE)
```

Advantages of designating data administrators. Designating data administrators leaves you free to have people administering Warehouse Manager security without requiring you to grant those people *ALLOBJ authority. If the user who will be in charge of implementing Warehouse Manager security does not have *ALLOBJ authority, it is a good idea to make that user a data administrator.

Disadvantages of designating data administrators. You cannot lock data administrators out of any data in any products (that is Query or Warehouse Builder) that use Warehouse Manager as their server. In other words, data administrators can see any objects for which they have iSeries authority.

Reviewing and Updating Your User Profile Scheme

In several of the recommended practices that follow, you will be encouraged to apply specific security settings to libraries and tables on a user-profile basis. This level of security administration can be tedious if you have many users. As an alternative, you may consider using group profiles. If the same security settings will apply to several users, for example all members of a given department, then it would be efficient for you to create a group profile and make those user profiles members of that group. You can then administer Warehouse Manager security for the entire group rather than for each individual member.

You can use Warehouse Manager Client to create group profiles and to make users members of the group.

Warehouse Manager products will honor authority if it is gained through an iSeries supplemental group. However, you cannot administer supplemental groups with Warehouse Manager Client.

Using Menu-Level Security

In a menu-level security scheme, you rely on a set of menus or application interfaces to restrict user access to underlying data. If an option is not on a menu, your users do not have access from an iSeries display session. Generally, in this scheme, your data files will not have object-level security applied to them; that is, you have not excluded specific users from specific files. Almost all client-server tools will bypass this security, because outside of the environment of the iSeries display session, you no longer have a menu interface protecting access to your files.

Menu-level security is considered a nonrestrictive security plan, and your goal when deploying STRATEGY should be to restrict users from gaining access to data that they must not see.

Option 1—To restrict all access and authorize users to only those files that they need:

- ▶ Using Warehouse Manager Client, change server options as follows:
 - Select to exclude users from libraries by default.
 - Select to exclude users from tables by default.
- ▶ Using Warehouse Manager Client, authorize specific users or groups of users to files they need to access.

Advantages. This option provides the most secure and least error-prone protection of sensitive data files. As libraries and tables are added to your database in the future, they are automatically protected from users until you specifically authorize them.

Disadvantages. The process of authorizing users correctly can be time consuming. This effort can be reduced if your data files are organized so that all files in a given library can be opened up to a user or group of users. If this is the case, you can set the server options to exclude users from libraries by default, without excluding them from tables by default. Then you need only to authorize a user to a library to allow him or her access to all tables in that library.

Option 2—To restrict access to specific files:

- ▶ Using Warehouse Manager Client, restrict specific users or groups of users from files that they should not be allowed to read.

Advantages. This method can be implemented quickly if you have only a few sensitive files and know exactly which files should not be viewed by which users.

Disadvantages. This method leaves a lot of room for error and is not recommended. If you don't remember or don't know about a sensitive table and fail to secure it, the data will be available to anyone. If sensitive tables are added to your database in the future, anyone will be able to read them until you restrict them.

Option 3—To run at security level 20:

Running your iSeries at security level 20 is a variation on menu-level security. At security level 20, all users have the special authority *ALLOBJ, meaning that they can access any information in your data files. If this is the case, you are probably relying on menu security to restrict access to those

files. However, outside your iSeries display session environments, users will not be restricted. If this is the case, you need to take additional measures to restrict access to your system.

- ▶ Using Warehouse Manager Client, change server options to Enforce ShowCase Security on *ALLOBJ profiles.
- ▶ Decide whether to use the approach described in the previous two options. If you want to restrict access globally and authorize users to specific files, continue by following the steps in option 1. If you want to allow access globally and restrict users from specific files, refer to option 2.

Advantages. This is the only way to allow Warehouse Manager security settings to take effect when the system security level is set to 20.

Caution: Menu-level security methods will protect you when your users are using STRATEGY products. However, other ShowCase ODBC products, IBM Client Access, or FTP will still let users access files and libraries. Thus, these applications will bypass your security scheme because they do not tie into Warehouse Manager's security, and there is no underlying object-level security.

Application-Level Security

With an application-level security scheme, all tables have object-level security applied to exclude users from accessing them. Applications then adopt authority in order to access necessary tables. When users attempt to use the tables with a client-server tool or any other application that is not designed to adopt the correct authority, they will still be excluded from the tables.

Application-level security is considered a restrictive security plan, and your goal when deploying STRATEGY should be to allow users to access the tables that they need but still be restricted from using the tables that they should not see.

Security on a New, Dedicated, Data Warehousing iSeries

If you decide to deploy security on a new iSeries that will be dedicated to your data warehousing applications, you will not need to consider an existing security scheme as the previous scenarios do. We recommend that you use a restrictive approach when deploying a new data warehousing system.

- Run the iSeries at system security level 30 or higher.
- Do not grant users *ALLOBJ authority unless it is necessary for system operations or other (non-STRATEGY) applications.
- Create all new database files with *PUBLIC authority set to *EXCLUDE.
- Use Warehouse Manager Client to authorize specific users to database files they will need when using Warehouse Manager products.

Advantages. This approach allows you to completely control Warehouse Manager security for STRATEGY users without the risk of users accessing database files that they should not see. If you use other applications on this iSeries in addition to your STRATEGY applications, you can set authorities for specific files needed by those applications without affecting your Warehouse Manager security.

Disadvantages. The flexibility of applying different security schemes for applications from different vendors means that you will need to administer these schemes separately.

Restricting Users to Libraries in their iSeries Library Lists

Library list security allows you to restrict users so that they can access only the set of libraries in their library list. Without library list security, users have access to all libraries.

To set library list security:

- ▶ Open Warehouse Manager Client and from the Manage menu, choose Server Options.
- ▶ Select the Library Access tab in the Server Options dialog box.
- ▶ Select Return only libraries in user portion of the library list.

Advantages. If you are currently using the library lists in job descriptions or system values to control user access to libraries as part of your system security policy, choosing library-list security allows you to continue using this mechanism for database file access. For example, the J.D. Edwards package uses library-list security in some cases.

Disadvantages. If library lists are not currently part of your database access scheme, you will probably find no advantage to using library-list security.

Applying Column Filters

Warehouse Manager column security allows users to see data in some columns of a database file, while restricting them from seeing data in other columns. For example, an employee master file might contain employees' names, titles, department numbers, hire dates, and social security numbers. Some of your users may need access to the information, but you want to keep the social security numbers confidential. Using Warehouse Manager column security, you can exclude all users or specific users from seeing the social security column, while allowing them to access all other columns.

To apply column security with Warehouse Manager Client:

- ▶ Ensure that the users have authority to see the database file and highlight the particular column that you want to protect.
- ▶ Choose a security option. If you want to restrict everyone from seeing the column, choose the user profile *PUBLIC. If you want to restrict specific users, choose those user profiles and specify the Exclude authority setting.

Advantages. Without column security, you would not be able to restrict users from any of the data in files that contain columns with sensitive information. Column security gives you much more flexibility in deciding which files your users may access.

Disadvantages. Warehouse Manager column security applies only when you are using the client products that use Warehouse Manager products.

Applying Row Security

Warehouse Manager row security allows users to see data in some rows of a database file while restricting them from seeing data in other rows. For example, a prospective customer file might contain prospect information for your company's potential clients around the world. However, you want only the sales managers in each country to be able to view information about prospects in their country. Using Warehouse Manager row security, you can exclude users from seeing the rows that contain information about prospects in other countries.

If Warehouse Manager row security is applied to a logical file (data view), that security will be used. If there is no security on the logical file, then security on the base physical file will be used. The field used for the row security constraint for the base physical file must also be in the logical file. Otherwise, the security will fail. This applies only to row security.

To apply row security with Warehouse Manager Client:

- ▶ In our example, you would ensure that each sales manager has access to the prospect file, highlight the name of the table, and choose row security.
- ▶ For the U.S. sales manager, select that person's user profile and, in the Row Security text box, enter `\COUNTRY='USA'` where COUNTRY is the name of the column containing country information. Then you would select, for example, the German sales manager and enter `\COUNTRY='DEU'` and repeat this for the sales manager in each country.

Another approach to row security is to write user exit programs. The PCSACC exit program on the CHGNETA command is supported. (However, *REGFAC is not supported and will be treated as *OBJAUT.) User exit programs specified on the PCSACC parameter can be used to essentially apply row security for all client-server programs that support this parameter.

Advantages. Row security gives you much more flexibility in deciding which files users may access. Row security enables you to restrict users access to data in files containing rows with information that they should not see.

Warehouse Manager Client gives you a way to specify this row security through a graphical user interface, without forcing you to write user exit programs. It also allows you to apply row security for STRATEGY applications without affecting your other applications.

Disadvantages. Warehouse Manager row security only takes effect when you are using the client products that use Warehouse Manager Server.

Setting Application Security

Warehouse Manager Client allows you to specify which users can run STRATEGY applications. Some STRATEGY applications allow users to do more than access files. For example, Warehouse Builder can create new files, and Warehouse Manager Client can apply security, alias, and resource settings. In most cases, you should exclude all users from using these applications except those specific users whose job it is to administer that application.

Authority to use Query and Report Writer can generally be given to all users. However, it is possible to create batch output files with these products. See the next section for tips on maintaining security with batch output.

System Security and STRATEGY Application Output

Warehouse Builder, Query, and Report Writer are all capable of creating new files on your iSeries. You should manage this capability carefully. Your goals when managing output from STRATEGY applications are:

- Make sure that new objects do not unintentionally expose data to users who were previously restricted from viewing it.
- Make sure that existing data is not accidentally overwritten.

Making Sure That Restricted Data Is Not Exposed

In some circumstances, it is possible for new objects to give users unintended access to data. For example, suppose you have set Warehouse Manager security to exclude some users from certain tables. But a user who does have authority to access that table might run a Query batch query and store the output in a file. If that new file is placed in a library to which the previously unauthorized users have access, they will now be able to view the data in that file. To prevent this situation with query output to a batch file, you should implement at least one of the following suggestions:

- Choose the server option to exclude users from libraries by default.
- Specify the name of the library to which a user can write batch output using the Resource Settings option for Query or Report Writer, and then ensure that library is secured from other users.

The same situation could happen with Warehouse Builder target files. It is recommended that you exclude users from access to these applications. Only the person who deploys data warehousing or administers database access should use Warehouse Builder.

Making Sure That Existing Data Is Not Overwritten

You can control the security required to clear the target table when running a Query or Report Writer query in batch mode. This is to prevent a user from accidentally replacing a production file to which they have authority to update.

There are two modes for handling batch output security:

- You must be the owner of the batch target table or have Manage authority to the target table before being able to replace the results.
- You must have system object-level security and at least Read authority to the target table before being able to replace the results.

To decide which mode to use, first examine the files on your system:

- If many of your production files have *ALL or *CHANGE authority for your query users, you should consider using the first mode to gain additional Warehouse Manager security. This way a user will not be able to accidentally replace production data with the results of a query.
- If you have appropriate object-level security, you may disable the additional Warehouse Manager security.

To change batch output security:

- ▶ Open Warehouse Manager Client and from the Manage menu, choose Server Options.
- ▶ Select the Batch tab in the Server Options dialog box.
- ▶ Select one of the following options to determine who is allowed to clear a result table:
 - The owner of the table or any user with Manage authority to the table.
 - Any users with iSeries *CHANGE authority to the table.

Essbase/400 Security

Essbase/400 security planning is a necessary part of building a multidimensional data warehouse. For detailed information on defining Essbase/400 security settings, see Chapter 6 or the *Hyperion Database Administrator's Guide*.

Controlling Query/Report Writer System Usage

One aspect of improving your overall system performance is to control how Query and Report Writer use system resources. This means that you must manage certain aspects of your iSeries system to control the impact your PC applications will have on your iSeries system. This can be done by setting maximum query run limits, setting server job run priorities, and controlling the level of parallelism on DB2 Symmetric Multiprocessing systems. These performance improvement methods are generally implemented by using Warehouse Manager Client.

Setting Maximum Query Run Limits

Using the query run limit, you can protect your system from long-running or ill-formed queries. Query run limit support is built on the predictive query governor support of OS/400. Using Warehouse Manager, you can set the maximum run length of a user's query. If the query optimizer on the iSeries predicts the query will take longer than the limit, it stops the query *before* it is run.

Advantages. Setting query run limits allows you to control the system resources that any query user is allowed to use. It also prevents system resources from being consumed by a single, ill-formed query.

Disadvantages. It limits the user to running queries that meet the criteria set by the database administrator.

Since the query optimizer is making an educated guess about the run time for a query, it may not be completely accurate (the query could actually have a shorter or longer run time than the estimate given).

Assigning Query Run Limits

To determine the query run limits, you should consider the types of queries that your users run and the appropriate protection that you need to provide for the overall system. To make this analysis:

- ▶ Determine the amount of time that it takes to run several typical queries. You can use the Performance Analyzer feature in Query to estimate the amount of time a query will take to complete, or you can use auditing in Warehouse Manager Client and run the Detailed Select Statement query.
- ▶ Determine the other requirements of your system and the job run priorities established earlier.
- ▶ Using this information, break your queries into two groups:
 - The set of queries that completes quickly enough to be run interactively.
 - The set of queries that should be run in batch.
- ▶ Use Warehouse Manager to set up the appropriate query limits for *PUBLIC, individual users, and groups of users.

Example

In this example, you want to protect the system from CPU-intensive queries. Use the Performance Analyzer in Query to estimate the run limit on a set of queries commonly used by a user named QUSER. Once the analysis is made, you will set up the following limits for user QUSER:

- For any queries estimated to run less than 5 minutes, run them interactively at priority 50.
- For queries estimated to take more than 5 minutes, run them in batch at priority 50.
- If the query is estimated to run more than 30 minutes, it will not be allowed to run.

Verifying the Run Limit

The query run limit is enforced by estimating how long a query will run. Since this is only an estimate, some queries that take longer than the limit will be allowed to run. Also, some queries may not be allowed to run even though they will finish within the specified time. To verify whether a query will run within the limits, use the Performance Analyzer in Query, which will inform you of the query's estimated run time.

Setting Server Job-Run Priorities

The job-run priority is used to prioritize the server job relative to other work on the system. The job-run priority is the standard iSeries work management run priority. The lower the number, the higher the run priority.

Advantages. Provides control of system resources by prioritizing jobs relative to other work on the system.

Disadvantages. Other work can starve the query if it has a higher priority.

Assigning Job-Run Priorities

- ▶ Consider all work performed on your system and rank it by importance.
- ▶ Based on ranking, choose the job-run priority that you want to use when running queries.
- ▶ Using Warehouse Manager, assign job-run priority for *PUBLIC.
- ▶ You may further refine priorities by assigning job-run priorities based on user ID or groups of users. This allows you to prioritize query usage among users.

Example

After evaluating your system, you must run the following jobs:

- Order-entry application
- End-user queries using STRATEGY
- Nightly batch, order processing application

Based on this, you assign the following job-run priorities:

- Order-entry application (priority 20)
- End-user queries using STRATEGY (priority 20-30)
- Nightly batch, order processing application (priority 50)

The order-entry application and batch processing application already had the run priorities defined. For queries, you defined a range of priorities so that you can further prioritize individual users as needed.

Using Warehouse Manager, you set up a general run priority of 20 for interactive queries and 30 for batch queries. By setting the value for *PUBLIC, you are applying it to all user profiles on the system. On some later date, you may define specific run priorities for individual users and groups of users. These will then override the *PUBLIC value.

Additional Information

Once the server is started, the job priority may be changed using the CHGJOB CL command. When a user connects to the iSeries or submits a query to batch, the following shows which run priority is used:

- Warehouse Manager setting for the user profile.
- Warehouse Manager setting for the group profile if the user is a member of a group.
- Warehouse Manager setting for *PUBLIC.
- The class description used by Warehouse Manager Server.

If the run priority does not exist at a particular level, the next level is used. In the previous example, you would use the value of *PUBLIC because there are no settings for the specific user profiles or the user's group profile.

For batch jobs, the Warehouse Manager setting is determined when the batch job is submitted, not when it is run.

Controlling Level of Parallelism

This section applies only if you have the DB2 Symmetric Multiprocessing for OS/400 feature on your iSeries. Controlling the level of parallelism will determine whether the system will allow a query to be logically split into multiple, independent tasks that can be carried out in parallel.

The iSeries system value QQRYDEGREE can be used to tune the level of parallelism when running queries on your system.

Assigning the Level of Parallelism

- ▶ Consider all work performed on your system and decide if queries can be allowed to consume additional CPU and I/O resources in order to finish quickly.
- ▶ Decide on the amount of parallelism that you are going to allow. Currently, the system value gives you the following control:
 - *NONE. Do not allow any parallelism. Queries are not the highest priority on the system and you cannot afford to allow the system to devote the resources necessary to an individual query.
 - *IO. Allow IO parallelism. You are willing to let the system use additional IO and mainstore resources to complete a query faster. In some cases, allowing the system this freedom to use more resources while running a query will improve the cumulative total resources used.
 - *OPTIMIZE. Allow the system to choose whether IO and CPU parallelism should be used to improve query performance and overall system performance. In some cases, allowing the system to use more resources while running a query will improve the cumulative total resources used.
 - *MAX. Allow IO and CPU parallelism. You want parallelism to be used whenever possible if it will help improve the individual query. You only use this system for running queries, so there is little impact on other types of applications.
- ▶ Once you have set the level of parallelism, review system impact and adjust it as necessary until you have the setting that works best for your environment.

Example

You have three systems in your company. You run queries on all systems, so you bought DB2 Symmetric Multiprocessing for OS/400 to get the additional data access methods and now want to use the appropriate level of parallelism. The following shows how each of your systems is used:

- SYS-A. You run an order-entry application on this system and run occasional queries.
- SYS-B. This is a development and test system with which you enhance your application and test your queries.
- SYS-C. This machine is a dedicated data warehouse system and is where most of your queries are run.

Based on analysis, you choose to set the QQRYDEGREE system value as follows for each of your systems.

- **SYS-A=*NONE.** Your order-entry application is your first priority. You cannot afford to have a single query use a majority of the system resources even for a short time.
- **SYS-B=*OPTIMIZE.** You let the system decide. You have no mission critical applications running on this system, and you want to test queries that sometimes use parallelism.
- **SYS-C=*MAX.** This system is used only for decision support. You want the queries to be fast, and you are willing to allow the query to use any resources necessary to complete.

Enabling Expert Cache

Expert cache is a set of improved paging algorithms used for database files and related objects. By enabling expert cache, you are improving query completion time and reducing system paging and faulting rates. Enabling expert cache should be the first step in tuning your system for running queries.

To Enable Expert Cache

- ▶ Determine the storage pool that the Warehouse Manager Server job is using. Normally, this will be *BASE (storage pool 2).
- ▶ On the iSeries, use the WRKSYSSTS command with the assistance level set to 2=Intermediate.
- ▶ Press F11 until you see the heading Paging Option. Enter *CALC over the top of the existing value to enable the expert cache algorithms. Press Enter.

You can also use the CHGSHRPOOL command.

Tuning System Performance

System performance tuning refers to adjusting pool sizes, time slice, and subsystem allocations to improve response time, throughput, paging rates, and contention.

System performance tuning can improve overall usage of system resources and individual query response times. However, in most cases, you must tune the *entire* system to realize the benefits. For this reason, we recommend the following:

- Always consider the entire system and all system work when adjusting system-wide resources.
- Never adjust system resources for the server without considering the rest of the system load.
- System tuning starts with system performance analysis. Do not tune the system without first completing this analysis.
- Tune system performance *after* following the other performance suggestions in this section.
- In most cases, it is best to hire a consultant who specializes in system tuning. System tuning can be an exercise in trial and error. A good consultant already knows what works and what does not in many situations.

If you decide that system-level performance tuning is right for your system, see the following sources for information to help you get started:

- *IBM OS/400 Work Management Guide*
- *IBM Performance Tools/400—Getting Started*
- An iSeries consultant specializing in system performance

DB2 Symmetric Multiprocessing for OS/400

The DB2 Symmetric Multiprocessing feature improves query performance by adding the following to DB2/400:

- Additional data access methods and optimizer improvements
- CPU parallelism for query processing
- Support for the QQRITMLMT and QQRIDEGREE system values

Warehouse Manager Server takes advantage of the functionality provided by this feature. And while its name indicates that it is for multiprocessor systems, it also provides improvements on single-processor machines.

Using Auditing to Understand Your Queries

STRATEGY can provide you with information about the impact STRATEGY users have on your system when accessing the iSeries relational databases using SQL. To collect this information, you must use Warehouse Manager Client and choose Auditing from the Manage menu. This information can help you make tuning and performance decisions about specific queries that your users are running.

Impact to Overall Performance

Turning on the auditing feature will have an impact on your system's performance, and this may be noticeable to end users. Auditing works by collecting information on every query run and writing that information to a database file; both the additional overhead of collecting information and storing output operations to disk can affect your performance. Because environments differ greatly, it is difficult to predict the exact degradation that you might see.

Because of the possibility of performance degradation, we suggest that you turn on auditing for a limited time only. Collect information for a period of time, turn auditing off, and analyze what you've collected. Do not leave the auditing feature permanently turned on. Rather, auditing should be used as part of a planned cycle of information collection and analysis.

Collecting the Information

- ▶ Using Warehouse Manager Client 3.0 or greater, connect to the iSeries data source (Warehouse Manager Server at release 3.0 or greater) in which you would like to audit your users' queries.

- ▶ From the Manage menu, choose Auditing.
- ▶ Set the options to collect data for a specified period of time or until the auditing output file reaches a specified size, and then click Start Auditing. If you decide to end auditing before the specified time or file size is reached, click Stop Auditing.

After you start auditing, information will be collected on SQL SELECT statements for each STRATEGY user who connects to the server library in which you started auditing using a ShowCase ODBC connection. These will probably be Query and Report Writer users. Local distributions run by Warehouse Builder will also be recorded. STRATEGY products released prior to release 3.0, including Data View Manager and Analyzer Data Modeler, will also be audited if they connect to a 3.0 server library in which auditing is started.

Note that some STRATEGY products, including Warehouse Builder and Warehouse Manager Client, use SQL statements internally to process user requests. In order to provide an accurate picture of the impact STRATEGY products have on your SQL usage, these statements are also audited.

In some cases, Warehouse Manager Server is able to optimize a user's Query or Report Writer query so iSeries SQL is not used. In these cases, you will not see that user's query recorded in the auditing file, because it did not affect SQL performance.

Analyzing the Information

After auditing, it is time to analyze the information. Because the analysis involves running SQL queries against the database of information collected, it is recommended that you begin analysis only after auditing has ended. Otherwise, your own queries on the auditing information will be added to the database file as you analyze it, and this may skew your view of your users' actual SQL usage.

In Warehouse Manager Client, you will find several queries designed to help you analyze this data. From the Reports menu, choose Audit Queries. Each query is provided in both a Query (.dbq) format and Report Writer (.rpt) format. To use these queries, you must have at least version 3.0 of Query or Report Writer. An explanation of each query follows, with some suggestions on how to use them.

Select Statements: Basic Information

This report returns basic information about the queries run while auditing was turned on. Use it when you simply want to see which users are running SQL SELECT statements and what the statements are.

Select Statements: Detailed Information

This report returns detailed information about the queries run while auditing was turned on. Use it when you want to start analysis of the SQL SELECT statements that have run on your system. You may choose to do further analysis after seeing detailed information on the STRATEGY products used, the time to execute statements, the query optimizer's access methods, etc.

For each query run, there is at least one record in the report per file used in the query. Each record shows which access method the DB2/400 optimizer chose to use when retrieving your data from

that file. Always note the job number and unique statement number of the query you're interested in so you can find information about that particular query in the other reports.

Ordering. For different analysis purposes, you may want to order the data in this report differently than the way it is presented. Because this report, in the Query format, uses Break Groups, there is a special way to do this. To order by different columns in the group, from the Format menu, choose Break Groups and drag the column by which you want to order to the top of the Break Groups list.

Timings. The detailed information report includes the run time for each query. The time reported is not "wall clock" time, it is an internal processor time as reported by DB2/400. This information is best used for comparisons between queries.

This time may not match exactly a time reported by Query or Report Writer performance analysis. Again, the auditing time is an internal database-reported time, while the performance analysis is an estimate of all processor time used by the application. All comparisons made between queries should be done with timing reports from the same tool, either auditing or performance analysis, but the results of the tools should not be mixed.

You may notice that the same SELECT statement has varying run times. The reason for this might be the system environment, including the load on the system at the time the query was run and the job priority given to the user. Another reason might be that the SELECT statement was run multiple times in the same job, and the product was able to reuse the statement without going through some of the overhead associated with its first run.

When users query views instead of files. Note that if users are directly querying a view or logical file, this report shows that view as the name of the file used. Similarly, other auditing queries count that view as the "file used" for the purpose of most frequently queried files. If you need to find which physical file is actually being used by a view or logical file, add columns QQPTLN and QQPTFN to your report.

Multiple records for a queried file. You might see more than one record for a file in this report because it might be joined to itself, so it actually was used more than once. Or the query might have been implemented with an encoded vector index (EVI), in which case the access plan is recorded as both an arrival and an existing index. You may use the Queries Which Used Existing Indexes query to verify whether an EVI was used. Another scenario is that the filename is actually a view or logical file that references multiple physical files.

Full Select Statement

This report returns the entire text of the SELECT statement. You may need to use this report if a SELECT statement is longer than 1024 characters. The auditing file breaks SELECT statements into several records if the text exceeds 1024 characters.

This report is most useful if you use the Report Writer (.rpt) version rather than the Query (.dbq) version because the .rpt allows you to see multiple lines of the statement text.

Queries That Build Temporary Indexes

Temporary index builds may be of interest to administrators because often they are time consuming. Use this report to find which queries the DB2/400 optimizer chose to build temporary indexes for, and whether the optimizer advises building a permanent index.

Queries That Used Existing Indexes

This report returns the list of all SQL queries that used an index already on the system. You may be interested in this information if you created indexes specifically for the purpose of improving a query's performance, and you want to verify the index was used.

You may notice for some queries that the index name and library are the same as the filename and library. In this case, it means the key on the physical file was used to run the query.

Indexes Advised

The DB2/400 optimizer sometimes advises you to build a permanent index to assist in query performance. This report tells you which queries had indexes suggested for them, what files the indexes should be built over, and what the key fields of the indexes should be. By building the suggested indexes, you may improve your users' query performance.

Most Frequently Queried Files

This report returns a list of the files most often specified in your users' SELECT statements. Use it when you want to know which are the most "popular" files to query. This may assist you in making information more readily available to users or in deciding how to optimize your environment.

Which Queries Use a Given File

If you see something in the previous report that makes you wonder exactly why a certain file is being queried so often, you may use this report to find out exactly which SELECT statements ran against that file and which users ran them. To use this report, you will need to enter the name of the file that you are interested in and its library.

Note: The information you enter in the prompts is case sensitive.

Query Governor Information

DB2/400's predictive query governor prevents queries from running if it estimates that they will take longer than a time value that you have specified. This report shows whether the query governor was active at the time a query was run and whether any users' queries were prevented from running because the predictive query governor estimated their processing time to be too long.

Run this report when you've used Warehouse Manager Client's Resource Settings to set Query Limits for a user. Setting Query Limits for a user enables the DB2/400 predictive query governor when that user connects to the server. This report can help you see whether you are preventing too many queries from running because the Query Limit is set too low.

Additional Analysis Tips

If you need more detailed information about a query than you can get from the auditing queries, you can cut the SELECT statement out of the report (this is easiest using Report Writer) and paste it into the SELECT statement edit box in either Query or Report Writer. Then choose Performance from the Run menu to get detailed messages about the access paths considered, temporary indexes or result sets built, or indexes used.

Managing Your Auditing Environment

Saving Auditing Results

When you start auditing, if you select the Result Data option to Overwrite existing file instead of Append to existing file, any previous data collected will be cleared. If you would like to save this previous data for reference at a later date, follow these steps before restarting auditing with the Overwrite option:

- ▶ Use the iSeries CRTDUPOBJ command to create a copy of the file *SCDBMON*, which you will find in the server library in which auditing started. Be sure to specify *YES for the Duplicate data option.
- ▶ In the same way, make a copy of the file *SCAUDINF*, also found in the server library.
- ▶ Make a note of the names of your new files and their libraries.
- ▶ Open the auditing query or report that you want to run against the copied data.
- ▶ From the File menu in Query or Report Writer, choose Save As and save a copy of the query or report.
- ▶ Using your new, saved copy of the query or report, choose Tables from the Query menu. For each reference to either *SCDBMON* or *SCAUDINF* in the query, click the table name, and then click the Change Link button. In the Change Link dialog box, enter the library name and the filename of the copy of *SCDBMON* or *SCAUDINF* that you made earlier. (You may leave any references to *SCAUDMAP* alone, because that file is neither cleared when you begin auditing nor updated when users run queries.)
- ▶ Repeat step 6 for each query or report that you need to run against the saved audit files.
- ▶ Remember to delete your copies of *SCDBMON* and *SCAUDINF* when you are finished with them. The copies that you save are no longer under STRATEGY control so Warehouse Manager Server will not automatically clean them up.

Subsequent Installations

If you reinstall Warehouse Manager Server in a library where auditing has been run (for example, to upgrade the library to the next release of Warehouse Manager Server), be aware that your auditing information isn't preserved during such an installation. If you want to save your data, you may follow the steps in the previous section, being sure to save the files in a library other than your server library.

Size of the Auditing File

The auditing file *SCDBMON* can become quite large. The record length of the file in V4R3 and V4R4 is 2958 bytes, and each SELECT statement can generate several records.

While you do have the ability to specify that auditing should end when the file reaches a certain size or when a given time period expires, be aware that any users who are currently connected when

these parameters are reached will continue to have their queries audited and information recorded in the auditing file until they disconnect.

Upgrading OS/400

Auditing should always be ended before you upgrade your version of OS/400. If you want to preserve any previous auditing information, follow the steps at the beginning of “Managing Your Auditing Environment” on p. 27. When you restart auditing after the upgrade, choose Overwrite existing file for the Result Data option.

Improving Query Performance

After improving your overall system performance by controlling Query and Report Writer system usage, it is possible to fine tune your queries by taking advantage of methods used by the iSeries to process queries and by following some general query creation tips.

Understanding the Query Optimizer

The iSeries uses a sophisticated piece of software called the **query optimizer** to efficiently retrieve end user and applications data. Understanding how the query optimizer processes queries will help you create queries that will run efficiently, since the query optimizer’s primary function is to turn an SQL statement into a set of structures that enable fast and efficient data retrieval from the iSeries. Like most optimizing software, query optimizer produces a query as follows:

- Breaks the SQL statements into functional “chunks.”
- Validates each chunk, for example, verifies the existence of the files and fields referenced in the query.
- Optimizes the query.
- Builds structures needed by database run-time code for data retrieval.
- Passes those structures to the database run-time component.

One of the key outputs of an optimized query is a control structure called an **access plan**. An access plan is a bookkeeping structure for tracking the objects used by DB2/400 and determining how the database functions will access and process each record in the database.

A key piece of information held by the access plan is the access path to be used by the query. The access path refers to the method(s) that the database management system (DBMS) used to perform the query. View the access plan as a road map and the access path as the road itself.

During the optimization phase, the query optimizer will evaluate many alternatives to implement a query. These alternatives may use a single access method or a combination of methods. After coming up with a list of possible access plans, the “cost” of each access plan is calculated based on the expected amount of CPU and I/O work each requires. Once the cost of each plan is available, the query optimizer chooses the cheapest method for processing the query.

For a more detailed explanation of how the query optimizer processes information, see the IBM document *DB2 for OS/400 Database Programming*.

Using Query's Performance Analyzer to Evaluate Your Queries

Improving query performance can often be an exercise in trial and error. You may need to analyze your query, make changes, then analyze it again until you have achieved an acceptable level of performance.

To tune your query:

- ▶ Understand how the query optimizer chose to implement the query.
- ▶ Analyze the run-time performance of the query and determine if it is I/O or CPU bound.
- ▶ Based on the results from steps 1 and 2, experiment with the techniques described in the remainder of this document to eliminate unnecessary CPU processing and disk I/O.
- ▶ Go back to step 1 until the performance doesn't improve or all possibilities have been attempted.

The most important step in improving query performance is to understand what choices the optimizer made for the query. Once this is understood, you can experiment and influence the optimizer's future decisions.

Query provides an easy way to analyze your query through its Performance Analyzer feature.

To analyze your query:

- ▶ After creating your query, from the Run menu, click Performance. The Performance dialog box appears.
- ▶ Click Analyze to analyze the efficiency of your query.
- ▶ Once the analysis is complete, click the Log tab to view the implementation messages for your query. The meaning of each of the optimizer messages can be found in the Performance Analyzer online Help.

To effectively use the detailed information provided by the log, see the following sources:

- Online Help. The online Help contains a description of the provided information along with some general performance-tuning tips.
- IBM document *DB2 for OS/400 Database Programming*
- IBM document *DB2 for OS/400 SQL Programming*

The Performance Analyzer feature of Query allows you to quickly and easily test queries before they run to determine how they were implemented and the estimated run time. This estimate is provided by the DB2/400 query optimizer and is fairly accurate. However, as mentioned previously, the optimizer doesn't always have adequate information for an accurate estimate. Note that Performance Analyzer is also available in Report Writer.

Creating Indexes to Improve Performance

Creating indexes can dramatically improve query performance. However, creating indexes will also slow the performance of OLTP applications (for example, order entry) because the index will be updated for every insert, update, and delete operation. Indexes built over large files will also consume additional DASD.

When should an index be built? It depends on the importance of query performance weighed against that of OLTP performance, unless you have separated queried data from updated data (a data mart or data warehouse environment). The rest of this section will focus on creating indexes to help query performance. The trade-offs should be considered before creating indexes for production files.

To determine what kind of indexes will help your query performance:

- ▶ Look at 5–10 of your most resource-intensive or frequently run queries. Using the performance analyzer, go through the query optimizer messages as described in the previous section.
- ▶ Identify the following attributes of these queries:
 - Join conditions
 - Group-by columns
 - WHERE clauses
- ▶ If the query messages gathered in step 1 point to join conditions or group-by clauses as the source of spontaneous index builds, seriously consider building permanent indexes above those files. If an index doesn't exist for a join condition, it is probably best to build an index by defining a join-logical file. Join-logical files not only supply the necessary index, but they eliminate the need for your end users to specify a join condition in their query.
- ▶ Examine your WHERE clauses. Pick out the most selective parts of the WHERE clause—that is, those that tend to eliminate the most records from the file. If those conditions are typically joined by AND with the rest of the WHERE criteria, those fields/columns are usually the best candidates for a permanent index. If multiple fields are usually joined by AND, the index should usually be built with those fields in the key. Multikey indexes should be ordered from most selective to least selective WHERE criteria.

The following documents provide more information on creating indexes:

- IBM document *DB2 for OS/400 SQL Reference*
- IBM document *DB2 for OS/400 SQL Programming*
- IBM document *DDS Reference*

The following examples present sample SQL statements and the types of indexes that might improve their performance:

Example 1

```
SELECT partnum, partname FROM orders WHERE orddate > 960530
```

The candidate for an index is `orddate`. The index should be created when queries referencing `orddate` eliminate 80% of all rows.

Example 2

```
SELECT custnum, partname FROM orders WHERE orddate > 960601 AND custnum = 976771
```

The candidate index would have key fields of `custnum` and `orddate`. The order of the key fields will depend on which is most selective (put the most selective first).

Example 3

```
SELECT ordnum, custname, partdesc FROM orders, parts WHERE orders.partnum = parts.partnum
```

The candidate index for `partnum` would be in either or both files. It could create a join-logical file instead of an index for even better performance.

Example 4

```
SELECT custname, sum(price) FROM orders GROUP BY custname
```

The candidate index is `custname`.

Example 5

```
SELECT custname, sum(price) FROM orders WHERE custnum = 976771 GROUP BY custname
```

The candidate index is `custnum` if the `WHERE` clause is very selective.

Note: To perform the group-by clause, the system will build an index at run time if one doesn't exist over `custname`. This will be fast when `custnum` is very selective because the existing index for `custnum` will contain few records. The index for `custnum` will be used to build the index for the group-by clause.

Changing ODBC.INI Settings

Several settings in the `ODBC.INI` file also influence the performance of Query. The first setting is `SQLOptimizationLevel`. It can be set to one of two values: 0 or 3. The default value is 3, which tells the `STRATEGY` components to optimize the query so that it fully utilizes all available APIs on the `iSeries` to minimize query run time. A value of 0 indicates that the query should be passed directly to the SQL components on the `iSeries`. There are some queries that will run better if the query is passed to SQL. While the recommended setting is 3, this is another technique that might help if other approaches don't work.

Another `.INI` setting is `AllowParameterization`. The default value for this option is `Yes`, which allows the ODBC driver to parse your SQL statements and to replace the literal values found in these statements with parameter markers. This parameterized statement is prepared and held for reuse during your Query session. If the statement is being run more than once during a Query session and only the literal values used in the statement are changed, the default option of `Yes` will help you avoid the constant SQL overhead at run time. If the majority of your statements are run once during a Query session and the literal values used in the statement aren't the only values changed, then `No` for this option will help the `iSeries Query Optimizer` make better estimates and decisions on your SQL, improving run time.

Query Creation Tips

Always Specify Join Conditions

If the query contains more than one table, it is almost always necessary to specify a join relationship between the files. If no relationship is defined, the system will perform a Cartesian product join that matches every row in one table with every row in another table. The typical result returns extremely long response times and large amounts of (nonsensical) data.

Limit the Use of Deeply Nested Subqueries

Subqueries complicate query optimization. As a result, the optimization phase has more opportunity to build an inefficient access plan. In general, it is a good idea to break the statement down into simpler queries that are run back-to-back if possible.

Limit Use of Data Type Conversion

Whenever possible, avoid forcing the system to do data-type conversions. For example, if you have a WHERE clause, such as WHERE INTEGER = 1.4, the system must convert the INTEGER field to another data type to accurately compare it to 1.4. The optimizer won't use any index built over INTEGER.

Avoid String Padding

In cases where you compare a string literal to a field, ensure the length of the literal matches the length of the field. If the literal is longer, the system must convert the field to the longer length (and pad with blanks) to do an accurate comparison. As with numeric conversion, this eliminates use of any index built over the character field. A shorter literal string won't cause problems, since the system will pad the literal appropriately to use the index.

Minimize the Use of LIKE

Avoid using pattern matching when possible. Particularly bad is the use of % or _ at the start of a pattern. Since the optimizer depends on the use of leading key fields in indexes, the use of a leading wild card character eliminates index use in these situations.

Avoid the Use of Arithmetic Expressions in WHERE Clauses

The use of arithmetic expressions often eliminates the use of an index that would otherwise apply to the selection criteria.

Optimizing Performance for Warehouse Builder

Following are several ways to improve the performance of your Warehouse Builder relational database distributions.

Choose the Most Efficient Type of Distribution

There are several types of distributions that Warehouse Builder provides. The various types of distributions are shown in Figure 3-1. The type of distribution that you choose will have a significant impact on the performance of your distribution and its affect on your iSeries system. The performance characteristics of the various types of distributions are described in Table 3-1.

Figure 3-1
Warehouse Builder distribution types

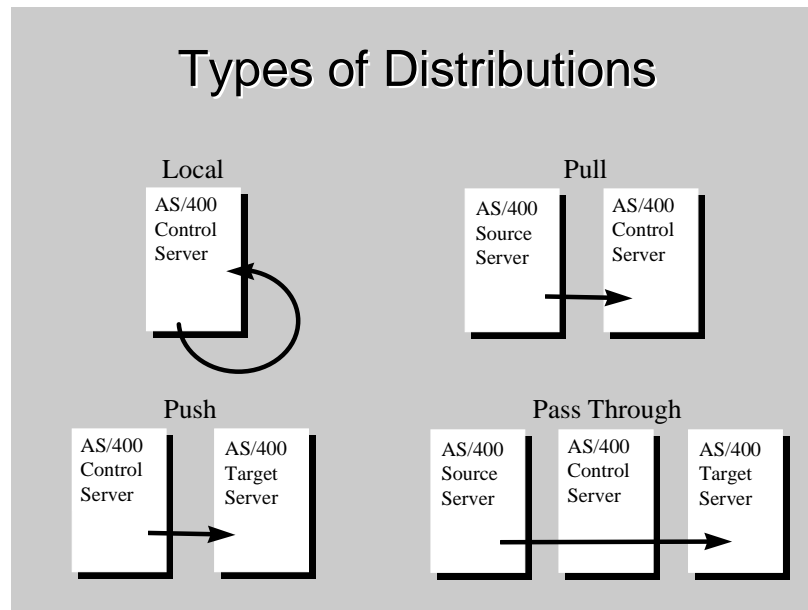


Table 3-1
Distribution types and performance levels

Type of distribution	Performance level
Local (within the same machine)	Very good
Pull (from a remote iSeries data source to a target table on the control server)	Very good
Push (from the control server to a remote iSeries)	Not very good
Pass through (from a remote iSeries to another remote iSeries)	Not very good

Use Pull Distributions

By far, the fastest type of distribution is a pull distribution. Pulling data from a target system is 10 to 15 times faster than pushing data to another system. This is due to the lack of blocking support in DRDA when doing inserts into a remote file. Therefore, we recommend that all performance-critical distributions pull data from the source system. This means the control server and the destination server should be the same system.

Creating Libraries

To store distributed data, Warehouse Builder will create libraries on the target system by default, improving distribution performance. To use journaling, you can specify that Warehouse Builder create collections instead. However, using collections and activating journaling may dramatically reduce distribution performance.

To create collections or libraries with Warehouse Builder:

- ▶ From the Tools menu, choose Options.
- ▶ Select the Default tab.
- ▶ Select Create Target SQL Collection. If this option isn't selected, Warehouse Builder may create libraries.

Using CLRPFM Command Instead of Mass Deletes

When the target table is defined in Warehouse Builder, you can specify that the target table is deleted on each distribution. There are two methods to accomplish this: SQL DELETE or CLRPFM (Clear Physical File Member).

When you define in Warehouse Builder to delete the target table, an SQL DELETE statement is issued to assure that the client code will work against any SQL data source. The SQL DELETE method is slower than CLRPFM but is applicable with all targets.

On the iSeries, the other method to remove all records from a file is to issue a CLRPFM command. CLRPFM is a faster method for deletion, but it is only valid with iSeries targets. To enhance performance, use the CLRPFM command before the data distribution is started. CLRPFM has two other distinct advantages:

- Removing old records is extremely fast and consumes few system resources.
- Disk space used by the deleted records is recovered.

To specify how to delete the target table:

- ▶ From the Tools menu, choose Options.
- ▶ Select the Default tab.
- ▶ In the Table Overwrite section, select the SQL DELETE or the Clear File Member Command option.

Sending Only Changed Records

Another way to improve distribution performance is to send only changed records between the source and target servers. While there is currently no automatic method to send only changed records between a source and target system (for example, as with products having journal-based replication services), the following paragraphs describe several methods to facilitate sending only changed records if you have a file getting only inserts and not getting updates or deletes.

Using CURDATE()

If there is a date in the file(s) referenced on a source system, you can use a variety of built-in SQL functions to determine what data are new since the last distribution.

For example, if you move data nightly from an orders file and that file has a date (either a date data type or a date that can be manufactured out of several fields and converted to a real date), then you can compare the date in the record with the current date (for example, `ORDERDATE = CURDATE()`).

Using a Sequence Number

If records on the source system have a field containing a sequence number (perhaps used as a key) that grows in value as new records are inserted, then you may be able to use this field to select only the recently inserted records. To use this method:

- ▶ Create a new file on the source system that has a field representing the `last_sequence_number`. Initialize this field to 0.
- ▶ When creating a definition, join it to the table created in step 1 and specify a `WHERE` clause to select only the records that are greater than `last_sequence_number`.
- ▶ Create a post-SQL statement that selects the maximum sequence number from the source tables and updates `last_sequence_number`.

This method assumes that no new records are being added to the source table between the time the main transfer executes and the time the post-SQL statement is executed to update `last_sequence_number`.

Other Performance Tuning Issues and Solutions

The efficiency of the `SELECT` statement running on the source system will have a direct influence on both the distribution time and the amount of system resources used. Therefore, when `SELECT` statements run against large files, it is recommended that you first use the Performance Analyzer feature of Query.

- Determine if the statement is optimized and/or how long, roughly, the query will run.

or

- Turn on auditing and run it as a local distribution or in Query to get the auditing information logged.

Based on this information, you may either adjust the SQL statement, build indexes to speed up the query, or perform the query in multiple passes (for example, divide the query into several smaller queries to reduce the complexity of the statement or perform more of the work on the target system instead of the source).

Managing a Multidimensional Database

This section focuses on performance tuning and planning a multidimensional database.

Optimizing Essbase/400 Performance

This section highlights important information for implementing Essbase/400. There are two important areas of optimization. The first is optimizing the Essbase/400 Server, and the second is optimizing the iSeries system on which the server is running.

Essbase/400 Performance Tuning

Issues critical to the success of your Essbase/400 implementation are load, calculation, and retrieval performance. Performance varies based on the relationships between the source data, the amount of data, the configuration of the Essbase/400 outline, the size and complexity of your retrievals, and the tuning applied to the database settings.

Note: It is important that you understand the concepts in the *Hyperion Essbase Database Administrator's Guide* and IBM's *iSeries Work Management Guide*.

Designing an Optimized SQL Data Load Plan

To optimize an Essbase/400 data load:

- ▶ Sort your input records so that the same sparse member combinations are grouped consecutively. In addition, sort your input records so that the sparse dimensions appear in the same order as the sparse dimensions in the outline. By sorting your data on the sparse dimension members, you minimize the number of times an Essbase data block must be read from disk, decompressed, enlarged, recompressed, and written back to disk. To sort your input data, specify the appropriate ORDER BY clause on the SQL statement in your rule file.
- ▶ Load the data from a data source on the server. If the data isn't located on the source, moving it to the server will improve load times since the data doesn't have to be transported over the network at load time.

For additional information and other optimization tips, refer to Volume II of the *Hyperion Essbase Database Administrator's Guide*.

Preventing Storage Overflow

Essbase data load or calc operations may exceed the predetermined Auxiliary Storage lower storage level limit and eventually crash the system. You can set up the iSeries to stop the data load and/or calc operation before storage overflow to prevent the system from crashing.

To stop the data load and/or calc operation before storage overflow:

- ▶ Set the QSTGLOWACN system value to *REGFAC.

- ▶ Register SCSTGEXIT to QIBM_QWC_QSTGLOWACN by typing the following at the iSeries command line:

```
ADDEXITPGM EXITPNT(QIBM_QWC_QSTGLOWACN) FORMAT(STGL0100) PGMNBR(1) PGM(ess-  
baselib/SCSTGEXIT)
```

where *essbaselib* is the name of the Essbase/400 installation library.

Note: The PGMNBR parameter must not already be in use.

- ▶ Once a system administrator has corrected the low storage situation (that is, once there is enough storage available), he or she must run the following command on an OS/400 command line before Essbase can resume normal functions:

```
RMVENVVAR ENVVAR(OS400STGLOWLMTREACHED) LEVEL(*SYS)
```

Optimizing the Essbase Calculation Process

Calculating an Essbase/400 model is CPU intensive and decisions that you make about the calculation process can affect performance. Refer to the *Hyperion Essbase Database Administrator's Guide* for more information about optimizing your Essbase calculations.

iSeries Performance Tuning

Performance can change depending on the system resources available to Essbase/400 (CPU, number and type of disk units, and main memory), the number of users accessing the data, and other concurrent activities on the system. For more information on tuning your iSeries, see the *IBM Work Management Guide*.

Once you understand your analyst's requirements, examine your implementation environment to determine the limits that will be placed on your solution. The following sections provide insight into issues to consider and questions to answer before continuing with the next phase of your implementation.

Balancing Priority between Warehouse Builder and Essbase

If you attempt to run Warehouse Builder distributions while running Essbase Server, be aware that these two activities compete for system resources. This is especially true if you are running on a single-processor system.

Depending on how you installed Warehouse Manager, Warehouse Builder distributions run a default priority of 20. If you start the Essbase/400 Server as a typical batch job, it runs at a default batch priority of 50. Both of these tasks are iSeries resource intensive. Since Warehouse Builder jobs run at a higher priority, your Essbase operations (load/calc/retrievals) will run very slowly.

Note: To prevent slow Essbase operations, set identical tuning priorities for your Essbase Server and data distributions. For example, if you run the Essbase Server at priority 50, set the data distribution to priority 50 as well. In addition, set your time slice under 25 milliseconds. This will prevent starvation of one of the jobs.

Controlling Memory Allocations

Essbase/400 performance is greatly impacted by the amount of memory it is allowed to use. Just like any other OS/400 application, you can manipulate the standard OS/400 work management techniques to control memory allocations for Essbase/400. On the iSeries, memory is allocated to various memory pools by OS/400. For example, memory can be allocated to the *BASE pool, *INTERACTIVE pool, or other user-defined pools. For more information on memory pools, consult the *IBM Work Management Guide*.

You can run Essbase/400 under the *BASE pool or under a dedicated memory pool. Use a memory pool that's sufficiently large and doesn't have excessive competition from other applications running on your iSeries. We suggest starting with a memory pool of 256MB for small databases accessed by a limited number of users. As you increase the number of databases, the size of the databases, or the number of users, consider increasing the amount of memory available for Essbase/400.

Because every Essbase/400 database is unique, we cannot specifically recommend system size and memory requirements. You can find a general Essbase requirements document on the Technical Support Web site. To access the document, go to <http://support.spss.com>. Once you log on, click Documentation and then click Technical Bulletins.

Grouping Databases into Applications

Typical Essbase/400 installations run many Essbase/400 databases on a single iSeries. While you can group several databases into a single application, a single application and database for each business problem is usually the best choice to simplify administrative tasks.

Essbase/400 and Symmetric Multiprocessing

Symmetric multiprocessing (SMP) refers to a system with more than one main processor. For example, some iSeries models can be equipped with two, four, eight, or more processors. In this case, SMP is a feature of the hardware. Multiple processors let the system simultaneously process more than one task.

On the iSeries, SMP can also refer to a DB2/400 software offering that can improve DB2 query performance regardless of whether the system has more than one main processor. This software offering gives OS/400 the enhanced ability to split the work to process a single query into several steps that can be run simultaneously in separate jobs.

Essbase/400 is implemented as a multithreaded program. Usually this allows Essbase/400 to take advantage of multiple processors.

Currently, the Essbase calculation process of a given database is single threaded and makes use of only a single processor. To use multiple processors during the calculation of a database, subdivide the database by creating transparent partitions. You can then take advantage of multiple processors by calculating each partition at the same time. For more information on transparent partitions, refer to the *Hyperion Essbase Database Administrator's Guide*.

Since each database is calculated under a separate thread, you can take advantage of multiple processors when calculating more than one database at the same time. Performance will vary depending on other iSeries resources such as the amount of available memory and the performance of your disk units. Performance also depends on other activity on the system.

If you have a single-processor system, limited memory, or limited I/O capacity (few disk arms), you will achieve better results by calculating only one database at a time.

Essbase/400 takes advantage of multiple processors while processing end user retrievals. This applies to retrievals submitted using the spreadsheet add ins, Analyzer, or user-written report scripts. For each retrieval, Essbase will use two threads. As multiple users run simultaneous retrievals, Essbase uses multiple threads to process these retrievals. Because Essbase is using multiple threads, it takes advantage of SMP and the work is spread across multiple processors.

Essbase Query Optimization Tips

This section contains tips for improving the performance of your Essbase queries. Many of these Essbase query optimization features can be found in the Essbase Options dialog box in Query and Report Writer.

Limiting Rows

If you have created a complex Essbase query or one that might return a large quantity of data, test the query first by limiting the number of rows returned. By doing so, the query will return only a specified number of rows, indicating that the query runs correctly. To limit rows, you can use any of the following three options:

- **Top/bottom sorting.** Use top/bottom sorting to return an exact number of top or bottom-ranked values in your report. For example, your database contains 2,000 cities, and you want to sort sales by city and display only the top 10 cities by sales. You can apply top/bottom sorting to one column per report. To apply top/bottom sorting, see “Applying top/bottom sorting” in the index of the Query or Report Writer online Help.
- **Suppressing rows that are missing all values.** Use this option to exclude rows containing only #MISSING values from your result data. #MISSING means the data in a given cell doesn’t exist. This may be because the member combination isn’t meaningful (for example, a given company doesn’t sell toboggans in Miami) or has never been entered. This option can be handy if you have a large number of rows with no values. To activate this option, from the Query menu, choose Essbase Options and then select Suppress Rows that are Missing All Values in the Essbase Options dialog box.
- **Suppressing rows that are all zeros.** Use this option to exclude rows containing only zeros from your result data. To activate this option, from the Query menu, choose Essbase Options and then select Suppress Rows that are All Zeros in the Essbase Options dialog box.

Suppressing Duplicate Shared Members

Suppress Duplicate Shared Members is an Essbase command that suppresses the display of duplicate shared members, reducing the amount of output data returned. Use this option to remove duplicate rows from your result data by excluding these duplicate shared members from your query. To activate this option, from the Query menu, choose Essbase Options and then select Suppress Duplicate Shared Members in the Essbase Options dialog box. For more information about shared members, see “Suppress Duplicate Shared Members” in the index of the Query or Report Writer online Help.

Creating Sparse Reports

Sparse reporting data retrieval is an Essbase feature that may improve data retrieval. Query and Report Writer can use this method to request data from the Essbase Server. The Generate a Sparse Report option uses the sparse data extraction method, which optimizes performance when a high proportion of the requested data rows are #MISSING. For more information on how sparse reporting works and when to use it, see “sparse reporting” in the index of the Query or Report Writer online Help.

Intelligent Sparse Reporting (ISR) is a separately licensed enhancement to Essbase/400 that enables faster query/retrieval processing for Essbase databases containing related dimensions comprised of a large number of members (defined as sparse dimensions to the Essbase engine). As an example, a database may contain several dimensions representing product SKUs such as product, color, and size.

For more information, contact your sales representative or choose Documentation from the STRATEGY menu.

Managing STRATEGY Data on the iSeries (AS/400)

Warehouse Manager Server supports STRATEGY PC applications. Data produced by STRATEGY applications is stored on the iSeries. You will need to perform the following tasks to protect this data:

- Move data from a test to a production environment
- Back up data
- Change physical files that have dependent data views
- Change data views that have dependent data views
- Manage passwords
- Manage information workflow

See Chapter 5 for more information on Warehouse Manager Server maintenance.

Moving Data from a Test to a Production Environment

Query, Report Writer, and Warehouse Builder give you the powerful advantage of developing queries, reports, data models, data views, and data distributions directly against iSeries or select NT or AIX databases. The environment in which you develop these files is critical to the success of daily database operations. If you develop these files in a production environment, you could adversely affect the performance of other operations. To prevent this problem, develop files against test data, either on a separate iSeries machine or in a separate library. Once testing is complete, you can put your queries, reports, data models, data views, and/or data distributions into production.

There are various ways to move data from a test environment to a production environment, as described in the following sections.

Saving and Restoring Data

This is a global approach, allowing you to back up all data residing on the iSeries. This approach is the most efficient, because you can move all information at once, but it limits you from moving individual queries, reports, data views, and so on. Use the Save/Restore procedure to save and back

up server information from one library and restore it to another. The target library can be on a different system or on the same system. See “Backing Up Data” on p. 43 for backup recommendations.

Redirecting Data

With this approach, you redirect a query, report, or data view from a test environment to run in a production environment. You do this on a file-by-file basis, which is helpful if you do not want to move all data at once through the Save/Restore procedure. However, this method requires more planning and structure. The key is to keep the structure of your libraries, tables, columns, and column data types consistent in both your test and production environments.

Sometimes, it’s necessary to modify the database structure by renaming libraries or tables or by removing or renaming columns. However, we recommend that you keep structural modifications to a minimum. If you are performing many structural modifications, it is wise to do a thorough database design review. Existing queries, for example, may cease to run after these kinds of changes. Change Link error messages are invoked when STRATEGY PC applications determine that libraries, tables, or columns are missing, or if column data types have changed. Consider the following two scenarios to help you decide how to structure your iSeries database:

Scenario 1: Testing and Production on Separate iSeries

If testing and production are performed on separate iSeries, maintain an identical database structure on both systems. That is, the library, table, and column names should match, and the column data types should be equivalent on both systems (for example, numbers to numbers or characters to characters).

	System	Library	Table	Column
From	SYS1	FY98	CUSTOMERS	CUSTNAME
	SYS1	FY98	ORDERS	SHIPDATE
To	SYS2	FY98	CUSTOMERS	CUSTNAME
	SYS2	FY98	ORDERS	SHIPDATE

After testing is complete, use the Change Data Source feature to switch a query, report, data view, data model, or data distribution to run on the production system. This feature is available in STRATEGY PC applications. Use the Performance Analyzer feature to evaluate the performance of a query on your production iSeries. This feature is available in Query and Report Writer.

Scenario 2: Testing and Production on the Same iSeries

If testing and production are performed on the same iSeries, maintain a set of test tables in a test library, parallel to the production library. The same concerns regarding matching names and data types apply here, except for the library names, which will be different.

	System	Library	Table	Column
From	SYS1	SAMPLE	CUSTOMERS	CUSTNAME
	SYS1	SAMPLE	ORDERS	SHIPDATE
To	SYS1	FY98	CUSTOMERS	CUSTNAME
	SYS1	FY98	ORDERS	SHIPDATE

After testing is complete, use the Change Link feature to convert a query, report, data view, data model, or data distribution to run against the production library. This feature is available in STRATEGY PC applications.

Creating Custom Library Lists for Users

Another approach for moving data from a test to a production environment is to create custom library lists for users. In this case, developers or testers run against a library list that points to test libraries; users run against a library list that points to production libraries.

To give users access to only those iSeries libraries in their library lists:

- ▶ Open Warehouse Manager Client and from the Manage menu, choose Server Options.
- ▶ In the Server Options dialog box, select the Library Access tab.
- ▶ Select Return only libraries in user portion of the library list.

This simplifies the database for users because all the data they need is contained in only a few libraries. However, it is difficult to troubleshoot database errors because tables are implicit in the libraries, thus making the exact location of data less obvious.

Backing Up Data

We recommend that you periodically back up Warehouse Manager Server. The purpose of the backup dictates how you back up the data. Will you be backing up temporarily so you can, for example, move server information between systems? Or will you be backing up as part of your disaster recovery plan?

Temporary Backups

The following situations require you to back up Warehouse Manager Server information temporarily:

- You need to install a new release of STRATEGY in the Warehouse Manager Server library.
- You need to delete the Warehouse Manager Server library but want to save the library's contents, such as security settings, queries, and data views.
- You need to move the Warehouse Manager Server library from a test system to a production system.
- You need to move Warehouse Manager Server contents from a test library to a production library on the same system.
- You need to distribute the data after you purchase a license for an additional iSeries.

For any of these situations, use the Save/Restore procedure to save and back up server information from one library and restore it to another. (See "After Restoring Server Information" on p. 46 for configuration issues that you will need to consider once Warehouse Manager Server information has been restored.) The target library can be on a different system or on the same system.

Using this procedure, you have the choice of saving server information to a file or a device such as a tape. In these situations, we recommend saving to a file. If you are moving data between systems, it is easier to send files electronically, especially for remote sites. If you cannot send files electronically between systems, you can always back up to a device.

Disaster Recovery Backups

To back up Warehouse Manager Server information for disaster recovery purposes, you should follow the same procedure as for other iSeries backups. Whether you use the iSeries GO SAVE command to back up the entire system or the SAVLIB command to back up individual libraries, you should always back up to a tape or other media device and then store the media off-site.

If you typically back up the entire system, Warehouse Manager Server information will be backed up when you perform the system save. For complete details about an entire system backup, refer to the *IBM OS/400 Backup and Recovery* documentation.

If you typically just back up individual libraries, simply include the Warehouse Manager library (for example, SCSEVER) in the list of libraries that you already back up.

To save your libraries:

- ▶ On an iSeries command line, enter SAVLIB and press F4.
- ▶ Enter the names of the libraries that you usually back up, plus the Warehouse Manager library name.
- ▶ In the Device prompt, enter the name of the device to save to (for example, TAP01 if you're saving to a tape).

If you need to restore information from a Warehouse Manager library that was saved in this way, we recommend that you restore this information from the device to a different library. That is, the

library to which you restore information should have a different name than the original library. If you restore information to the original library, that library could become corrupt, and you will lose data.

To restore the Warehouse Manager library:

- ▶ On an iSeries command line, enter RSTLIB and press F4.
- ▶ At the Saved Library prompt, enter the name of the Warehouse Manager library to be restored.
- ▶ At the Device prompt, enter the name of the device from which to restore this library.
- ▶ Press F9 to display more prompts and then page down to the next screen.
- ▶ In the Restore to Library prompt, enter the name of a different library (for example, SCLIB) in which you want to restore Warehouse Manager Server information. This library should not be the original Warehouse Manager library.

After restoring the original Warehouse Manager library (for example, SCSEVER) to another library (for example, SCLIB), install a fresh copy of STRATEGY in the original library. The installation procedure is described in the *STRATEGY Installation Guide*.

To recap, you have restored the SCSEVER library from a tape back to a library called SCLIB, and then reinstalled STRATEGY in a new SCSEVER library. (From this point forward, the original Warehouse Manager library is referred to as SCSEVER and the other library as SCLIB.)

Next, we recommend saving the server information in the SCLIB library and restoring it using the Save/Restore procedure. That is, you will want to run the SAVSRVRINF command from the SCSEVER library against the SCLIB library and then run the RSTSRVRINF command from the SCSEVER library to restore the information.

See Chapter 5 for detailed steps of the Save/Restore procedure. An abbreviated version of this procedure, as it relates to disaster-recovery backups, is described below:

- ▶ Add the SCSEVER library to your library list. On an iSeries command line, enter ADDLIBLE SCSEVER and press Enter.
- ▶ Save the server information from the SCLIB library. On an iSeries command line, type SAVSRVRINF and press F4. This command will run from the new SCSEVER library that you installed. Enter the following:
 - Enter SCLIB for the Warehouse Manager Server library.
 - Enter *SAVF to save the data to a file. A file is the most convenient choice, because you will keep the server information “live” instead of archiving it to tape.
 - Enter the save filename and the name of the library in which you want the file to be created. Please note that the file will be created if it does not already exist.

- ▶ Press Enter to save the information.
- ▶ Restore the server information to the SCSEVER library. On an iSeries command line, enter RSTSRVRINF and press F4. This command will run from the new SCSEVER library that you installed. Enter the following:
 - Enter *SAVF to restore the data from a save file.
 - Enter the name of the save file from which to restore and the name of the library in which you created that file.
- ▶ Press Enter to restore the information.

Backing Up Enterprise Reporting and Analyzer OLAP Data

The following command will save the server information including Analyzer OLAP functionality and any Enterprise Reporting information:

```
SAVSRVRINF SRVRLIB(serverlib) DEVICE(*SAVF) SAVF(savelib/savefilename) INFTYPE(*inftype)  
TGTRLS(*CURRENT) DLTVIEW(*NO)
```

where:

- *serverlib* is the library in which Warehouse Manager Server is installed.
- *savelib* is the name of the library in which the backup file will be stored.
- *savefilename* is the name of the backup file.
- **inftype* is the information type value.

To back up Enterprise Reporting data, use *AOS for the *inftype value. For general backup procedures, we recommend using *ALL as the *inftype value.

After Restoring Server Information

After running RSTSRVRINF to restore Warehouse Manager Server information, you may need to further configure your server so that it will run correctly. If you restored Warehouse Builder information, consider the following:

- If you restored the information to an iSeries different from the one on which it was saved, you may want to change your distribution sets and definitions. For example, if you were performing “pulls” of data from one iSeries to a test system, and you moved the distribution information to your production system, change the relational database (RDB) name of the target system so that the “pull” will now pull data to your production system. To make these changes, from the Tools menu in Warehouse Builder, choose Definition Properties.
- If you restored to a different iSeries, you should also ensure that TCP/IP communications are configured and active between the new system and the systems defined as source and target servers for the distributions.

- If you restored to a new iSeries or to a new library on the same iSeries, any schedule information defined for your distribution sets and any scheduled queries/reports are not restored. From the Sets menu in Warehouse Builder, choose Work with Warehouse Builder Sets to edit your sets and set up schedules.
- If you restored Warehouse Builder information into a server library already containing distribution definitions, the object IDs of the distribution sets will change. If you are running the sets from your own CL programs with the object IDs hard-coded in either the STRDD or SBMDD command, the hard-coded references may need to be changed. To avoid this problem in the future, use the object name of the distribution set in the STRDD or SBMDD commands. See “Running Distribution Sets with STRDD and SBMDD” on p. 50.

If you restored Warehouse Manager administration or Query information, consider the following:

- You can display the physical file RSTSRVPF to see the actions performed on the server catalog files.
- If you restored Warehouse Manager administration or Query information to a different iSeries, or if your environment has changed substantially since setting up this information, clean up the catalogs. From the Manage menu in the Warehouse Manager PC application, choose Clean Up ShowCase Catalogs.

Changing Physical Files That Have Dependent Data Views

Query lets you create and manage simplified views of any iSeries database. This simplified view, or data view, is saved on the iSeries as an SQL view. Since SQL views are saved on the iSeries and reference physical files (also known as tables), you may not be able to change physical files that have dependent views. Query includes several predefined queries and reports that you can run to find out which data views are built over a given iSeries table. You can run these predefined queries and reports from the Tools menu in Query.

If you find there are data views dependent on a physical file that you need to change, you can use the CHGPF command, which allows you to change a physical file without deleting dependent views. As a precaution and for recovery purposes, back up your views before using CHGPF.

To do this, use the Save/Restore procedure. For more information see “Saving and Restoring Data” on p. 41. That is, you will want to run the SAVSRVRINF command on the iSeries to save data views. In this case, leave Delete Views After Saving set to *NO when running the SAVSRVRINF command.

Changing Data Views That Have Dependent Data Views

Just as data views can reference physical files, data views can also reference other data views. If you built a data view over another view, you cannot delete, modify, or replace the base view because it has a dependent view. For example, if VIEW2 is built over VIEW1, you cannot delete, modify, or replace VIEW1. You have three options, all of which can be done within Query:

- Delete the dependent view, VIEW2, if it's no longer needed.
- Edit the dependent view to remove the dependency.

- Create a copy of the base view, VIEW1, and give the copied view a unique name (for example, VIEW1TEMP). Change the dependent view, VIEW2, to use the new view, VIEW1TEMP. This removes the dependency on VIEW1. You are now free to make your changes to VIEW1. After completing the changes to VIEW1, link VIEW2 back to VIEW1. You can then delete the temporary view VIEW1TEMP.

For complete details about data view tasks, see the Query online Help.

Managing Passwords

Scheduled reports and Warehouse Builder distributions rely on consistent passwords to connect to the iSeries. When a profile tries to connect to an iSeries through a STRATEGY application, the password is passed through to logon to the iSeries. In time, changed and/or expired passwords can cause the STRATEGY control tables to have incorrect passwords for a user. This may cause Warehouse Builder distributions and scheduled reporting to fail.

Changing Passwords with a STRATEGY Product

If the user's iSeries password has changed or expired, the user can update the applicable STRATEGY control table settings with the SCCHKPWD command. The SCCHKPWD command searches the local system and lists and/or updates the necessary control table settings with the new password. This ensures uninterrupted scheduled reports and Warehouse Builder distributions.

Updating Passwords with the SCCHKPWD Command

To update an iSeries user profile's password in the STRATEGY control tables, use the SCCHKPWD command. For the command details, press F1 after entering SCCHKPWD on the command line or see Appendix A. SCCHKPWD allows you to:

- Specify whether to update the tables or to generate a list of the items matching the search criteria.
- Limit the generated list to matching entries for DB2 databases, Warehouse Manager Servers, and non>ShowCase ODBC data sources.

Examples

- Update passwords for an iSeries RDB (Warehouse Builder distributions and Enterprise Reporting ODBC STRATEGY data sources):

```
SCCHKPWD USERID(myuserprofile) OLDPWD(MYOLDPWD) UPDATE(*YES) RDB(iSeriesname) AOS(*NONE)
ODBC(*NONE) NEWPWD(MYNEWPWD) PWDCONFIRM(MYNEWPWD)
```

- Update passwords for a non-STRATEGY data source (ER scheduled reports and WB Oracle or SQL Server pulls):

```
SCCHKPWD USERID(myuserprofile) OLDPWD(MYOLDPWD) UPDATE(*YES) RDB(*NONE) AOS(*NONE)
ODBC(mysource) NEWPWD(MYNEWPWD) PWDCONFIRM(MYNEWPWD)
```

- Update passwords for the local system (RDB and AOS):

```
SCCHKPWD USERID(myuserprofile) OLDPWD(MYOLDPWD) UPDATE(*YES) RDB(*LOCAL) AOS(*LOCAL)
ODBC(*NONE) NEWPWD(MYNEWPWD) PWDCONFIRM(MYNEWPWD)
```

- Get a list of items affected by a user password change:

```
SCCHKPWD USERID(myuserprofile) OLDPWD(MYOLDPWD) UPDATE(*NO) RDB(*ALL) AOS(*ALL)
ODBC(*ALL)
```

- Get a summary of RDB, AOS, and ODBC names affected by a user password change:

```
SCCHKPWD USERID(myuserprofile) OLDPWD(MYOLDPWD) UPDATE(*NO) DETAILS(*NO) RDB(*ALL)
AOS(*ALL) ODBC(*ALL)
```

Using SCCHKPWD with Essbase Build Plans

Using the SCCHKPWD command, you can update SQL passwords (for DRDA connection) that are used with an Essbase build plan.

Note that the DRDA system name is not stored in the control table (that information is stored in the Essbase rules file and read by Essbase during actual data load or dimension build). To update the SQL password, use the associated Essbase Server to identify the correct record. Be aware that because the same Essbase Server can use different DRDA data sources for data load using the same user ID and password combination, you may update all of those passwords at once.

Example

One user has Essbase Server SCEssbase and uses:

- UserA and pwd123 for data load from DRDA data source SYS1
- UserA and pwd234 to load from SYS2
- UserA and pwd123 to load from SYS3

Running the following command will update user A's password for *both* SYS1 and SYS3:

```
SCCHKPWD USRID(userA) OLDPWD(pwd123) UPDATE(*YES) DETAILS(*NO) RDB(*none) AOS(*none)
ODBC(*none) ESSBASE(SCEssbase) NEWPWD(newPwd) CONFNEWPWD(newPwd)
```

The result may be undesirable if the user only intended to update the SYS1 password.

There are two ways to avoid this problem:

- ▶ Change user A's password on SYS3.
- ▶ Use different/unique user ID for each DRDA system connection.

Managing Information Workflow

Running Distribution Sets with STRDD and SBMDD

The STRDD (start data distribution) and SBMDD (submit data distribution) commands run Warehouse Builder distribution sets by set name (OBJNAME) or ID (OBJID). For more information on these commands, see Appendix A.

For example, the following commands run the distribution set Sales:

```
STRDD OBJNAME('Sales') OBJID(*NONE) or STRDD OBJNAME('Sales')
```

The following commands run the distribution set that has an ID 51:

```
STRDD OBJNAME(*OBJID) OBJID(51) or STRDD OBJID(51)
```

The SBMDD command works in the same way.

The STRDD and SBMDD commands can be inserted into a user program to run distribution set(s). These commands return an RSC6336 status message when a distribution set fails. The user program can monitor for the RSC6336 message, and the appropriate response can be coded into the user program.

For example, if one set in the program fails, the user program can stop instead of proceeding to the next set. This enhancement does not affect existing user programs that use STRDD or SBMDD. This is a status message and the program is not required to monitor for it.

Example

Program using STRDD:

```
PGM /* PGM1 */
ADDLIB LIB(SCSERVER)
STRDD OBJNAME('Sales')
MONMSG MSGID(RSC6336) EXEC(DO)
    SNDMSG 'Set 1 Failed, call application support - Sets 2
    and 3 not run'
    GOTO END
ENDDO
SBMJOB CALL(PGM2)
SBMJOB CALL(PGM3)
END:
ENDPGM
```

Pulling iSeries Data into Essbase NT

It is possible to submit your Essbase NT load plans from STRATEGY Warehouse Builder to Essbase on the NT (6.0 patch 1 or later) platform. Once created, you can use all the features of Essbase on the NT platform.

Requirements

Be sure to meet these requirements before you begin this procedure.

- The SCODBC ShowCase ODBC driver must be installed on the Essbase NT Server. Use the STRATEGY installation CD to install one of the following STRATEGY PC Client components: Query, Warehouse Builder, Warehouse Manager, Analyzer, or Analyzer Designer.
- Set up a ShowCase ODBC data source on the Essbase NT (6.0 patch 1 or later) Server. From the STRATEGY menu, select Microsoft ODBC Administrator.
- Set up Hyperion Essbase SQL Interface on Essbase NT (6.0 patch 1 or later). This is a separately licensed add-on offered by Hyperion. You can verify licensed products in the Essbase Application Manager by clicking Server | Information | License Info.
- In Warehouse Manager, you must enable STRATEGY licensing for this feature by adding a password for the Essbase NT application. This password is available by contacting a service representative. You can also add this password with the CHGSCAUT command on the iSeries. See “Add Password command” in the Warehouse Manager online Help.
- Start your Warehouse Manager Server and Essbase NT Servers.

Creating Application and Database, Database Outline, and Data Load Rule Files

Before you start, make sure your Warehouse Manager Server and Essbase NT Servers are running.

Note: For instructions on creating the new application and database, saving the database outline to Essbase NT Server, and saving the data load rule files to the Essbase NT Server, see the Hyperion Essbase documentation.

Verify and Edit the Data Load Rule Files on Essbase NT

After the data load rules are saved to Essbase NT Server:

- ▶ Highlight the Application Desktop Window for Essbase NT Server. Select the application and database that you just created.
- ▶ Click Data Load Rules.
- ▶ Select a data load rule file and click Open. The data load rule file will appear.
- ▶ From the File menu, choose Open SQL. The Selected Application, Server, and Database dialog box will appear. Click OK. The Define SQL dialog box will appear.
- ▶ In the SQL Data Sources list box, verify that the selected data source points to the correct SCODBC data source.
- ▶ In the Select text box, make any necessary changes to the SQL Select Statement, which determines which data you're extracting from the iSeries data warehouse.
- ▶ In the From list box, verify that the library listed (the format is library.table) is the library on the iSeries containing the data files.

- ▶ In the Define SQL dialog box, click OK/Retrieve to retrieve sample data. The iSeries data will appear.
- ▶ Finish creating and editing the rule file. For more information, see the *Hyperion Essbase Database Administrator's Guide*.
 - You may now use Application Manager to perform a dimension build/data load.
 - In Warehouse Builder Client, you can reference this load rule in your Essbase build plan. When defining the Essbase build plan, type the Essbase NT IP address.

STRATEGY Integration with J. D. Edwards (JDE)

STRATEGY 4.5 supports versions A7.x and A8.x of JDE World and version B7.33.x of JDE OneWorld. The following sections describe enhancements to STRATEGY integration with JDE for this release. For additional JDE information such as data dictionary and security information, go to <http://support.spss.com> and log on to the ShowCase division support site.

Non-iSeries Profile Sign-On

Non-iSeries profiles cannot be used with Warehouse Builder or Warehouse Manager Client, and they cannot be administrated with Warehouse Manager Client. Also, non-iSeries profiles are not supported through Enterprise Reporting (ER). This means you cannot sign on to ER via a browser with non-iSeries profiles, and you cannot save to, open from, or log on to the Server in Query or Report Writer with a non-iSeries profile.

Previously, JDE OneWorld users needed an OS/400 user profile for STRATEGY to recognize their environment information. Administrators can now use the SETJDEOWA command to enable STRATEGY to recognize environment information for OneWorld users who do not have an iSeries user profile.

- ▶ To enable SETJDEOWA, use the following syntax:

```
SETJDEOWA SINGLUSRPRF(*ENABLE) JDEOWUSER(owprofile) JDEOWPASS(owpassword)
```

where *owprofile* is the user's JDE OneWorld profile and *owpass* is the user's JDE OneWorld password.

Note: The job description for the iSeries profile used in the SETJDEOWA command must have the JDEOW, QTEMP, and QGPL library lists. The profile must also have an initial program of BV3C in QGPL in order for the JDE OneWorld application user to be recognized as a OneWorld user. For more information about iSeries user profiles and job descriptions, refer to your iSeries documentation.

- ▶ To disable SETJDEOWA, use the following syntax:

```
SETJDEOWA SINGLUSRPRF(*DISABLE)
```


Editing ShowCase.cfg to Use an Alternate F0094 File

Usually, the F0094 file in JDE OneWorld contains the Master Library List data. JDE OneWorld integration with STRATEGY relies on the data in this file to set the library list upon connecting to the database through the Web interface or the application interface. On occasion, the JDE OneWorld environment administrator may choose to omit the F0094 file in its OneWorld client environment and use information contained in a client access data source. Since STRATEGY cannot determine the data source used for the library lists, a method has been provided so that users can manually specify an alternate file (which must contain F0094 data). This method involves duplicating the F0094 file and editing the configuration file (*ShowCase.cfg*).

To set up this alternate configuration:

- Find the JDE system library name by using the WRKOBJ command:

```
WRKOBJ OBJ(*ALL/F980WSEC) OBJTYPE(*FILE)
```

The object and library name will be displayed.

- A duplicate of the F0094 file is required. The CRTDUPOBJ command may be used, with the *F0094SPSS* filename in *MYLIB* and the JDE system library name *SYSB733* as examples:

```
CRTDUPOBJ OBJ(F0094) FROMLIB(SYSB733) OBJTYPE(*FILE) TOLIB(MYLIB) NEWOBJ(F0094SPSS)
DATA(*YES)
```

Note: The DATA(*YES) option duplicates the data. If you want to populate an empty file with the library list information, you can set the option to DATA(*NO).

- The main configuration file (*ShowCase.cfg*) in the Warehouse Manager Server library must contain two entries specifying the location of the alternate filename and library name. The iSeries EDTF command may be used to edit the file.

Client access may be used for the same purpose. Following is the EDTF example, where *SCSERVER* is the Warehouse Manager Server library name:

```
EDTF '/SCSERVER/proddata/config/ShowCase.cfg'
```

Add the following two lines to the configuration file:

```
showcasecorp.JDEOneWorldF0094File=F0094SPSS
```

```
showcasecorp.JDEOneWorldF0094Library=MYLIB
```

- To activate this change, stop and restart Warehouse Manager Server. For more information, see Chapter 2.

User-Defined Codes (UDC)

During the initial setup of the JDE system, an administrator will customize descriptions for users based on category codes that exist with the JDE system. A user can select these UDC customized descriptions in the STRATEGY columns display. In STRATEGY 4.5, UDC customized descriptions are available in Query, Report Writer, and when running interactive reports via a browser using Enterprise Reporting.

Note: To work with Warehouse Builder definitions that contain JDE data and UDC columns, OS/400 V4R5 users must apply the following PTF or the definitions will not run: 5769SS1 V4R5M0 SF64377.

Multiple Library List Environments

STRATEGY retrieves library list environments that have been set up within the JDE environment. A user who has been set up in the JDE system to have multiple library list environments will have the option to select an environment on the initial connection or have one set as the default. In STRATEGY 4.5, the prompting for library list environments can take place when running interactive reports via a browser using Enterprise Reporting. An enhancement has also been made allowing the prompting for library list environments if the environment is defined for the user group instead of the individual profile.

Selecting Group Library Lists

There is an order that JDE must follow to select group library lists. This order differs for OneWorld and World users.

OneWorld

The order of library list selection for OneWorld users is:

- User with a single library list (taken from the *F0092* file)
- Group with a single library list (taken from the *F0092* file)
- User with a multiple library list (taken from the *F0093* and *F0094* files)
- Group with a multiple library list (taken from the *F0093* and *F0094* files)

World

Currently, JDE World depends on the initial programs—J98INIT or J98INITA.

The order of library list selection for World users with J98INIT is:

- User with a single library list (taken from the *F0092* file)
- Group with a single library list (taken from the *F0092* file)

The order of library list selection for World users with J98INITA is:

- User with a multiple library list (taken from the *F0093* and *F0094* files)
- Group with a multiple library list (taken from the *F0093* and *F0094* files)
- User with a single library list (taken from the *F0092* file)
- Group with a single library list (taken from the *F0092* file)

Note: If the library list cannot be set due to invalid data, missing libraries, or the lack of a library list, Query, Report Writer, and Enterprise Reporting will not function properly.

Warehouse Manager Server Maintenance

This chapter discusses typical maintenance procedures that system administrators must perform on Warehouse Manager Server. For instructions on maintaining the server for Analyzer OLAP functions, see “Analyzer OLAP Maintenance” on p. 75.

Using TCP/IP with STRATEGY

STRATEGY software is used in a transmission control protocol/Internet protocol (TCP/IP) network environment. STRATEGY client and server products communicate with each other using connections through a TCP port.

Note: Whether you installed Warehouse Manager Server with one port number or two port numbers, you can use the CHGSCSRVA command to configure the port(s).

To change the TCP port number or port numbers after the install, use the CHGSCSVRA command. This makes it easier to migrate from a test environment to a production environment. See “Migrating Users to a New Warehouse Manager Library” on p. 58.

Configuring Warehouse Manager TCP/IP with STRSCSVR

If you configured Warehouse Manager TCP/IP during installation, you can skip this section. Before configuring TCP/IP, verify that OS/400 TCP/IP is configured and active. To enable Warehouse Manager TCP/IP support, run the CHGSCSRVA command.

The following additional commands are available in the server library to help users maintain Warehouse Manager Server TCP/IP support:

- CHGSCSVRA to change subsystem information and TCP/IP-specific objects.
- STRSCSVR to start the manager job.
- ENDSCSVR to end the manager and client jobs (to which the client is talking).

Finding Available Port Numbers

A TCP/IP port number is a numeric alias for an application. You assign a specific port to a Warehouse Manager Server installation library so the client software knows to which TCP/IP server application to connect. To enable the server for TCP/IP, you must assign a unique port number for each library. To view a list of the ports that are in use, run the OS/400 command WRKSRVTBLE.

The command CHGSCSVRA assigns the port number 43419 by default. If this number already appears when you run the command WRKSRVTBLE, choose another port number. The next available number should work.

Migrating Users to a New Warehouse Manager Library

At times, it may be necessary to have your STRATEGY users change libraries. For example, you may install a new release of the Warehouse Manager Server and first verify that it works in your test environment. You then want users to stop using the old release and begin using the new one.

In a TCP/IP network, users' client configuration employs the TCP/IP port to determine which Warehouse Manager installation it connects to. This allows you to stop the old server, configure the new library with the port number from the old library, and start the new server. This way, users can automatically connect to the new library without any changes to their client configuration.

Checking Usage of a Warehouse Manager Library

Listing Connections to Relational Applications Only

The following procedures list any TCP/IP active jobs to STRATEGY relational products such as Warehouse Builder. To list all the active connections to STRATEGY applications including Analyzer OLAP applications, see "Listing All Connections to All STRATEGY Applications" on p. 59.

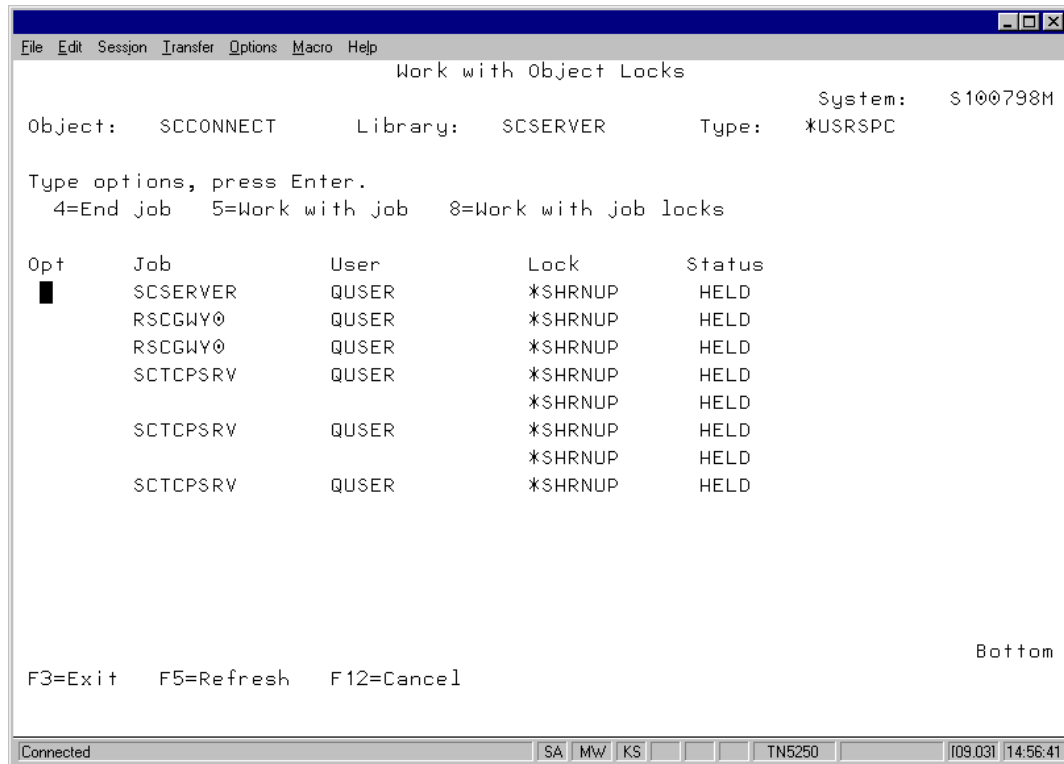
For Warehouse Manager libraries, run the following command:

```
WRKOBJLCK OBJ(serverlib/SCCONNECT) OBJTYPE(*USRSPC)
```

where *serverlib* is the name of the library that you want to check.

A screen similar to the following will appear. If TCP/IP is active, the screen shows the TCP/IP Manager job for the library (job name SCSEVER, which is the same as the library name) and may show TCP/IP connections (job name SCTCPSRV).

Figure 5-1
Using the WRKOBJLCK command to check library usage



To see which users are connected (for example, to ask them to disconnect before an installation), use Work With Job Locks (option 8). You should see a lock on the user profile (enter *USRPRF). (If you see QUSER, this indicates no user is connected). From the user profile name, you should be able to determine who is using the STRATEGY product. Alternatively, you can use the Warehouse Manager Client software to view active connections.

Listing All Connections to All STRATEGY Applications

This method lists active connections to all STRATEGY applications, including those using Analyzer OLAP functionality.

- ▶ Open the Warehouse Manager Client application and connect to the data source.
- ▶ Double-click the data source for which you want to see connections.
- ▶ From the Warehouse Manager menu, choose Applications.
- ▶ From the Manage menu, choose Active Connections.

Ending STRATEGY TCP/IP (Stopping the Server)

Before ending TCP/IP, verify that no one is using the library (Refer to “Checking Usage of a Warehouse Manager Library” on p. 58). If you see the library name in the list of jobs with locks, or jobs named SCTCPSRV are listed, TCP/IP is enabled.

From the library in which you want to end TCP/IP, run the command ENDSCSVR. This command will end TCP/IP clients and stop the TCP/IP manager job for the library.

Uninstalling a Warehouse Manager Library

You may occasionally need to uninstall Warehouse Manager libraries on your iSeries system, such as test libraries.

Uninstalling Warehouse Manager Server Libraries

- ▶ Log on as QSECOFR or with equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).
- ▶ Back up any Warehouse Manager settings that you want to save. See “Backing Up Warehouse Manager Server Information” on p. 61 for instructions.
- ▶ Verify that no users are connected to the library. See “Checking Usage of a Warehouse Manager Library” on p. 58.
- ▶ Restore the UNINSTSC utility from the UNINSTSC save file that exists in the Warehouse Manager Server library and enter the following command:

```
RSTOBJ OBJ(*ALL) SAVLIB(QTEMP) DEV(*SAVF) SAVF(serverlib/UNINSTSC) MBROPT(*ALL) ALWOBJ-  
JDIF(*ALL) RSTLIB(QSYS)
```

where *serverlib* is the name of the Warehouse Manager library to be uninstalled.

Note: This will restore three objects into QSYS.

- ▶ Interactively enter the Warehouse Manager command UNINSTSC and then press Enter. This command cannot be submitted to batch.
- ▶ Enter the name of the Warehouse Manager library that you want to uninstall and then press Enter to confirm.

Uninstalling Warehouse Manager Server Libraries (pre-STRATEGY)

- ▶ Log on as QSECOFR or with equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).
- ▶ Back up any Warehouse Manager settings that you want to save. See “Backing Up Warehouse Manager Server Information” below for instructions.
- ▶ Restore the UNINSTSC utility from the UNINSTSC save file that exists in the Warehouse Manager Server library. This must be restored from a 3.0 or later release.

```
RSTOBJ OBJ(*ALL) SAVLIB(QTEMP) DEV(*SAVF) SAVF(serverlib/UNINSTSC) MBROPT(*ALL) ALWOBJ-
JDIF(*ALL) RSTLIB(QSYS)
```

where *serverlib* is the name of the Warehouse Manager library to be uninstalled.

- ▶ Interactively enter the Warehouse Manager command UNINSTSC. It cannot be submitted to batch.
- ▶ Enter the name of the Warehouse Manager Server library that you want to uninstall.
- ▶ To uninstall the library, press Enter.

Backing Up Warehouse Manager Server Information

Because of the way Warehouse Manager is installed, we recommend backing up your installation library with SAVSRVRINF, rather than using conventional means. Following is the series of procedures that you should use to back up a Warehouse Manager Server library.

Saving STRATEGY iSeries Media and Passwords

Be sure to save the iSeries media and passwords that you received. You will need them to restore the server library. These passwords will be valid when restoring to the same iSeries.

Saving Server Information

Note: If you are saving information and the device is tape, STRATEGY objects must be the first or the only objects on the tape. Otherwise, the save to a tape will fail. Also, before saving Server information, be sure to end the Warehouse Manager Server (for instructions, see Chapter 2).

This task is required only if you are using any of the following applications:

- Warehouse Manager (server options such as iSeries library access and default public authority to iSeries libraries and tables; security and resource settings, aliases, and column attributes; and application security).
- Query (data views).

- Enterprise Reporting.
- Warehouse Builder (such as data distribution sets and definitions).
- Analyzer applications (such as forms, models, pinboards, and views).

The Warehouse Manager, Query, Warehouse Builder, and Report Writer applications store information on the Warehouse Manager Server on the iSeries. Analyzer applications store information such as views and pinboards. Save and back up the information on the server with the SAVSRVRINF command, located in the server library, and then restore using the RSTSRVRINF command.

The Save and Restore feature, which consists of these two commands, is also useful for moving data from one server library to another. See “Moving Data from a Test to a Production Environment” on p. 41 and “Backing Up Data” on p. 43 (both in Chapter 4) for recommendations on moving data from a test to a production environment and backing up data.

To back up information that STRATEGY applications store on the Warehouse Manager Server:

- ▶ Log on with a user ID that has *ALLOBJ authority.
- ▶ Add the Warehouse Manager library to your iSeries library list and ensure that you have only one Warehouse Manager library in your library list. Run the command:

```
ADDLIBLE serverlib
```

where *serverlib* is the name of the Warehouse Manager library with which you want to work.

- ▶ Enter the command SAVSRVRINF and press F4.
- ▶ At the Library prompt, enter the name of the server library from which server information should be saved. This library is typically SCSEVER.
- ▶ At the Device prompt, enter the name of the device to save to, or type *SAVF to save to a file. Press Enter. If you entered *SAVF, enter the save filename and the iSeries library in which you want to create the save file.

Note: Don't forget to stop the Warehouse Manager Server (for instructions, see Chapter 2).

- ▶ At the Type of Information to Save prompt, enter one of the following:
 - *ALL to save all of the information described in the following bullets.
 - *ADMIN to save only Warehouse Manager administration information (such as security settings and aliases).
 - *AOS to save only Analyzer OLAP information (such as forms, models, pinboards, reports, and views) and Information Delivery applications' user information (such as shared queries).
 - *DIST to save only Warehouse Builder information (such as data distribution sets and definitions).
 - *DATAVIEW to save only Query data views.

Note: Even if you save all information, you can restore individual groups of settings.

- ▶ Press F10 to see additional prompts for the SAVSRVRINF command. These prompts do not need to be changed under most circumstances. However, the prompts are described below, in case they might be useful to you:
 - At the Target Release prompt, enter *PRV if you intend to restore saved information on the previous release of the OS/400, or leave the default of *CURRENT if you intend to restore saved information on the same or a later release of the OS/400.
 - At the Delete Views After Saving prompt, specify whether or not to delete data views after saving. It is recommended that you leave this option set to *NO. However, if you need to change physical files that have dependent views, enter *YES to delete the data views on your system. (The views can later be restored with the RSTSRVRINF command.) Another option is to use the CHGPF command, which allows you to change a physical file without deleting dependent views.

- ▶ Press Enter to save the server information.

Restoring Warehouse Manager Server Information

Installing STRATEGY iSeries Media and Passwords

- ▶ Follow the procedures in the “Installing Server Software” chapter of the *STRATEGY Installation Guide*.
- ▶ Enter your Warehouse Manager passwords. For more information, see the *STRATEGY Installation Guide*.

Restoring Server Information

Note: Before saving server information, be sure to end the Warehouse Manager Server (for instructions, see Chapter 2).

To restore the server information that was saved with the SAVSRVRINF command:

- ▶ Distribute your save file or move your tape or other device to the target system, if it’s different from the first system.
- ▶ Add the Warehouse Manager library to your iSeries library list and ensure you only have one Warehouse Manager library in your library list. Run the command:

ADDLIBLE serverlib

where *serverlib* is the name of the Warehouse Manager library with which you want to work.
- ▶ Enter the command RSTSRVRINF and press F4.

- ▶ At the Device prompt, enter the name of the device from which to restore or enter *SAVF to restore from a save file. Press Enter. If you entered *SAVF, enter the save filename and the iSeries library in which you created the save file.
- ▶ Choose whether or not to restore Warehouse Manager Server administration information. Press Enter. If *YES, the Run Catalog Check Program prompt displays. Choose one of the following values for this prompt:
 - *NO to not run the catalog check program on the restored information.
 - *AUTODELETE to delete any records that do not apply to objects found on the system.
 - *LISTONLY to simply create a list of records that do not apply to objects found on the system. You can later delete these records by running the CHKSRVRCAT command and specifying *YES on the DLTORPHAN parameter.

Note: To view the results of a catalog cleanup operation, use the Warehouse Manager Client application. On the Manage menu, choose the Clean Up ShowCase Catalogs command and then click View Results to display the results.

- ▶ Choose whether or not to restore Analyzer OLAP information. Press Enter. If *YES, the Duplicate User Option and Restore Configuration File prompts display. Choose one of the following values for the Duplicate User Option prompt:
 - *KEEP to keep the existing AOS data in the target user-permission file. That is, existing system manager privileges and user settings will not be replaced with any duplicate AOS data being restored.
 - *REPLACE to replace the existing AOS data in the target user-permission file with any duplicate AOS data being restored.

Choose one of the following values for the Restore Configuration File prompt:

- *YES to replace the existing AOS configuration file with the saved (SAVSRVRINF) version of this file.
 - *NO to not restore this file and thus not overwrite the existing AOS configuration file.
- ▶ If you specified *YES for either the Restore Warehouse Manager Information or Restore AOS Information prompt, the Duplicate Record Option prompt is displayed. This prompt specifies the action that restore takes if it finds any saved Warehouse Manager or AOS data that duplicates data in the target server library. Choose *KEEP or *REPLACE.
 - ▶ Choose whether or not to restore Warehouse Builder information. Press Enter. If *YES, the Restore Logs prompt is displayed. Choose whether or not to restore the logging and message information from previous runs of the data distribution sets and definitions that you saved.
 - ▶ Choosing to restore Warehouse Builder information also displays a duplicate record option. Choose whether to keep or replace duplicate records by entering *KEEP or *REPLACE.
 - ▶ Choose whether or not to restore Data View information.

- ▶ You can press F10 to see additional prompts for the RSTSRVRINF command. The following prompt does not need to be changed under most circumstances. However, it is described below in case it might be useful to you.
 - In the Catalog File to be Processed prompt, the default *ALL restores all catalog files that contain STRATEGY administration information. These files contain STRATEGY application resource settings, Warehouse Manager security settings, ShowCase aliases, and column attributes. Alternately, you can specify a single catalog file to restore. Press F1 on this prompt to display the Help, which lists the catalog files and their contents.
- ▶ Press Enter to restore the server information. Several status messages will flash and then a screen with additional configuration steps will appear. Read this information carefully, especially if you restored Warehouse Builder, Warehouse Manager, Query, or Data View information.

For example, if you are restoring Warehouse Builder information to a new server library on the same system or to a different system, we recommend you use Warehouse Builder to make any necessary changes, to create target tables, and to reschedule jobs. Alternatively, you can use the Warehouse Builder post- save/restore utility found on the product's Tools menu.

Detecting and Removing Obsolete Information

When using Warehouse Manager's Alias Manager and Security Manager, you define alias names or security rules. Alias names reference existing objects on your system. If those objects change, the alias names may reference objects or fields that no longer exist. Similarly, Security Manager defines access rules for a particular object and user. In this case, not only can the referenced objects change but the users may no longer exist.

When you use Query to create data views, Warehouse Manager stores information about these data views. It is possible that these data views could be renamed or deleted by programs other than Query. The Warehouse Manager information could then reference objects that no longer exist.

Given these situations, Warehouse Manager provides a way to resynchronize the catalog information with your changed system. On the Manage menu, click Clean Up ShowCase Catalogs. The command allows you to determine and view information that is no longer valid, and remove it if you choose. For additional information on the command, see the Warehouse Manager Client online Help.

Restricting Access to a Warehouse Manager Library

It may be necessary to put a Warehouse Manager Server library in a restricted state. This restricted state prevents users from connecting to the server library. Once a server library is in a restricted state, you can easily remove the restricted state so users can reconnect.

Applying the Restricted State

To put a server library in a restricted state:

- ▶ Log on as QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).

- ▶ Add the Warehouse Manager library to your iSeries library list and ensure that you have only one Warehouse Manager library in your library list. Run the command:

```
ADDLIBLE serverlib
```

where *serverlib* is the name of the Warehouse Manager library with which you want to work.

- ▶ Change the server status by entering the command CHGSRVRSTS and pressing F4.
- ▶ At the Library prompt, enter the name of the Warehouse Manager library that you want to put in a restricted state. This library is typically SCSEVER.
- ▶ At the Restrict Server Library prompt, enter *YES to prevent users from connecting to the server library. Only your current job (running the CHGSRVRSTS command) and the TCP/IP master job (if it's running) remain connected to the restricted server library.

Note: If users are already connected to the server library, it will not be put in a restricted state unless you force the restriction.

- ▶ Press Enter.
- ▶ At the Force Server Restricted State prompt, choose whether or not to force the server library into a restricted state:
 - *NO to attempt to put the server library in a restricted state. However, if users are connected to the server library, the attempt will fail.
 - *YES to force all connections to the server library to end (except for those listed in the fifth step—your current job and TCP/IP master job). This will end the jobs of any users who are currently connected to the server library.
- ▶ Press Enter. If you specified *NO in the previous step, the server library may or may not be put in a restricted state. If you specified *YES, the Maximum Force Wait Time prompt displays. This prompt allows you to set the maximum time that the CHGSRVRSTS command should run while waiting for all connections to the server library to end. The default time interval is 30 seconds. After you specify a time limit, press Enter.

The countdown begins *after* all connections to the server library have started to close. The server will be put in a restricted state once either of the following occurs:

- No more connections are found during the specified time interval.

or

- The time interval expires and there are no more connections.

If the time expires and there are still connections to the server library, the server will not be put in a restricted state. In this case, you will need to retry until the restricted state is achieved so that you can change the Enforce AS/400 Security option.

Note: Regardless of whether or not the server library is put in a restricted state, when the CHGSRVRSTS command is issued with the Force=*YES option, all active connections will begin to close. This process cannot be reversed once it has started. If the server library fails to enter a restricted state in the time interval specified, users whose connections were ended in the process will be allowed to reconnect, as will any other users.

Removing the Restricted State

To release a server library from a restricted state:

- ▶ Log on as QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).
- ▶ Add the Warehouse Manager library to your iSeries library list and ensure you only have one Warehouse Manager library in your library list. Run the command:

ADDLIBLE serverlib

where *serverlib* is the name of the Warehouse Manager library with which you want to work.
- ▶ Change the server status by entering the command CHGSRVRSTS and pressing F4.
- ▶ At the Library prompt, enter the name of the server library that you want to release from a restricted state.
- ▶ At the Restrict Server Library prompt, enter *NO to allow users to access the server library again.
- ▶ Press Enter to remove the restricted state.

Working with Journals

Warehouse Manager creates two journals when it is installed. These journals are used to maintain data integrity when catalog files are updated. The two journals are RSCSPCTLJ, which is associated with stored procedure catalogs, and RSCSCCTLJ, which is associated with Warehouse Manager, Warehouse Builder, and Query Client catalogs.

The main concern with journaled files is the space consumed by the journal receivers. To ensure the journal receivers do not consume too much space on your system, you can use one of the following approaches:

- Monitor the size of the receivers and issue a CHGJRN command when they get too large. Once the CHGJRN command is issued and new journal receivers are attached to the journal, you can (optionally) save the receivers and then delete them.
- Let the system maintain the journal receivers. Enter the commands in the following steps to do so:

```
CRTJRNRCV JRNRCV(serverlib/receiver-name) THRESHOLD(500000)
```

```
CHGJRN JRN(serverlib/RSCSPCTLJ) JRNRCV(serverlib/receiver-name) SEQOPT(*RESET) MNGRCV(*SYSTEM) DLTRCV(*YES)
```

```
CRTJRNRCV JRNRCV(serverlib/receiver-name) THRESHOLD(500000)
```

```
CHGJRN JRN(serverlib/RSCSCCTLJ) JRNRCV(serverlib/receiver-name) SEQOPT(*RESET) MNGRCV(*SYSTEM) DLTRCV(*YES)
```

Warehouse Manager and iSeries Work Management

This section describes how Warehouse Manager is influenced by iSeries work management controls and provides information for adjusting these controls. Before reading this section, you should be familiar with iSeries work management concepts. For information on iSeries work management see the *IBM Work Management* document.

Work Management Controls

Following is an overview of what you can control with each of the work management constructs as they relate to Warehouse Manager and running queries. Included is information for changing the constructs to influence Warehouse Manager behavior.

In general, you should not need to adjust the work management behavior of Warehouse Manager. This information gives you a basic understanding of how Warehouse Manager is interacting with your system and provides the information necessary for advanced systems management, such as system tuning or configuring job classes to control temporary storage limits.

Job Class

The job class used by Warehouse Manager can be used to control the following attributes with minimum impact to other jobs. Other attributes can also be controlled through the class description, such as time slice.

- Job run priority
- Maximum temporary storage
- Default wait time

Job run priority can be used to control Warehouse Manager's run priority relative to other jobs on the system. If this is the only attribute you want to control, using the Resource Settings feature of Warehouse Manager Client is an easier alternative.

Maximum temporary storage can be used to control the amount of storage a job uses at run time. Temporary storage is the storage required for Warehouse Manager and system programs at run time, internal system objects used by the system for this job, and temporary objects used when running a query.

Default wait time can be used to control the maximum time, in seconds, that the Warehouse Manager job will allow for an instruction to finish running. Typically, this will be the time Warehouse Manager waits to obtain either an object or record lock while processing a query.

Use the default wait time if locks are common in your environment and are affecting your query. Reducing the wait time will allow Warehouse Manager to return record and object lock errors more quickly. Lengthening the wait time will allow Warehouse Manager additional time for the object to become available.

In most environments a wait time of 30 seconds to 2 minutes is recommended. A wait time of *NOMAX is allowed but is not recommended as lock conditions may cause the PC to appear stopped for long periods of time.

To create a job class, use the CRTCLS command. Add or update the routing entries to use the new job class.

Storage Pools

The **storage pool** on the iSeries is a portion of the main memory isolated for a set of jobs. All jobs using a storage pool compete for the resource but do not compete against jobs using a different storage pool. Storage pools can be used to:

- Isolate memory-intensive jobs from other jobs on the system
- Give particular jobs more memory by reducing contention
- Keep highly used objects in main memory

Use caution when creating storage pools since the memory will be subdivided and will not be used unless a job is actively using the storage pool.

More than one storage pool can be associated with a subsystem. The routing entry indicates the storage pool that should be used by the job when running in the subsystem.

Subsystems

The subsystem Warehouse Manager runs in can be used to control the storage pools available to the server and to configure job classes within the subsystem.

To ensure the storage pool you are using has expert cache enabled, see “Enabling Expert Cache” on p. 22 in Chapter 3.

The subsystem the server job uses is determined as follows.

TCP/IP. TCP/IP uses the subsystem specified during the Warehouse Manager installation or with the CHGSCSVRA command. This will be either QCMN or QBASE by default on most systems. If you need different PC workstations in different subsystems using TCP/IP, you must have multiple Warehouse Manager libraries installed.

Batch queries. Batch queries use the subsystem servicing the job queue. This job queue was specified in the job description and is used when the query is scheduled by the user.

Routing Entries

Routing entries allow you to choose the job class used by Warehouse Manager and the storage pool within the subsystem the job uses.

TCP/IP jobs. TCP/IP jobs use the job class associated with the routing entry added when STRATEGY was installed or when the CHGSCSVRA command was issued. By default, STRATEGY uses the QWCPCSUP job class for TCP/IP connections.

To change the job class used by TCP/IP, use the CHGRTGE command. The routing data format of the TCP/IP routing data is listed as follows:

Table 5-1
TCP/IP routing data format

Data	Position	Length
Library	1	10
SCTCPSRV	10	8

The iSeries also manages TCP/IP. These jobs simply start the associated STRATEGY client jobs when a connection is made to the iSeries using ShowCase ODBC. The routing data is as follows:

Table 5-2
iSeries TCP manager job routing data

Data	Position	Length
Library	1	10
SCTCPMGR	10	8

Batch queries. For batch queries, you can add your own routing data in the job description. When that job description is a specified batch job schedule, the routing data will be passed to the system. You can set up routing entries based on the routing data you define to assign the appropriate job class.

Job Descriptions

The job description allows you to control the initial library list used by the server, job log logging options, and job accounting.

TCP/IP jobs. For TCP/IP, the job description of QUSER is used initially. The library list of the job description specified in the actual user's user profile then replaces the library list of QUSER. The job log logging and job accounting are not updated.

Batch queries. For batch queries, it's the job description specified when scheduling the batch job.

User Profiles

The user profile allows you to control:

- Object level security and management attributes applied when running the server.
- Job description for the user.
- Permanent storage limits.
- National Language Support attributes of the job.
- Output queue used when producing a job log or other spooled output.

Warehouse Manager will always run using the user profile specified at ODBC logon when connecting to the iSeries.

When using WRKACTJOB, WRKUSRJOB, or a similar command, the job name may contain the actual user profile or it may contain QUSER, depending on how the job was started. QUSER is the user profile many STRATEGY jobs run under. QUSER must not have a storage limit. To make sure that QUSER has no limit, run the following command:

```
CHGUSRPRF USRPRF(QUSER) MAXSTG(*NOMAX)
```

To identify a job by user, you can use one of the following techniques. To list all jobs by user, use the command:

```
WRKOBJLCK OBJ(myuser) OBJTYPE(*USRPRF)
```

where *myuser* is the user profile for which you want to list jobs.

To list all jobs for a particular library, use the command:

```
WRKOBJLCK OBJ(libraryname/SCCONNECT) OBJTYPE(*USRSPC)
```

Security Issues

Command Authorities

This section lists the provided commands and their authorities as set at installation.

On installation, existing command authority will be preserved for all users except *PUBLIC. *PUBLIC will unconditionally be set to *EXCLUDE or *USE as described in the following tables.

Existing Data Administrators will not have authority to access the commands. You may:

- Grant authority to those users by employing the iSeries security management commands.
- Use Warehouse Manager Client to grant the Data Administrator privilege.

The Warehouse Manager-provided iSeries commands in the following table are shipped with *PUBLIC authority set to *EXCLUDE. A security officer can grant appropriate authority using the standard OS/400 security management commands.

Table 5-3

Commands shipped as *PUBLIC = *EXCLUDE

Command	Description
CHGSCSVRA	Configure TCP/IP connectivity support, subsystems, and various Analyzer parameters
APYSCPTF	Apply STRATEGY PTF
CFGDD	Configure data distribution
CHGAOSA	Use CHGSCSVRA command
CHGSCAUT	Change STRATEGY product authority
CHGSP	Change stored procedure
CHGSRVRDFT	Change server default attributes
CHGSRVRSTS	Change server status
CHKSP	Check stored procedure
CHKSRVRCAT	Check server catalog information
CPYSP	Copy stored procedure
CMPSCJAR	Compile the STRATEGY Java files

Command	Description
CRTSP	Create stored procedure
DLTSP	Delete stored procedure
DSPSP	Display stored procedure
ENDSCSVR	End TCP/IP connectivity support
RMVSCAUT	Remove STRATEGY authority
RSTSMPLB	Restore STRATEGY sample database
RSTSP	Restore stored procedure
RSTSRVRINF	Restore server information
SAVSP	Save stored procedure
SAVSRVRINF	Save server information
SETLIBACC	Set library access
STRSCSVR	Start TCP/IP connectivity support
UNINSTSC	Uninstall server

The commands in Table 5-4 are shipped with *PUBLIC authority set to *USE.

Table 5-4

Commands shipped as *PUBLIC = *USE

Command	Description
DSPAOSINF	Display AOS information
DSPSCINF	Display STRATEGY tips
DSPSCAUT	Display STRATEGY product authority
SBMDD	Submit Warehouse Builder distribution set
STRDD	Start Warehouse Builder distribution set
STRRPTEEXEC	Start STRATEGY job executor

When granting a user Data Administrator privilege with Warehouse Manager Client, that user must be given iSeries *CHANGE authority to access the Warehouse Manager commands in the following table. If the user is not given *CHANGE authority, some administration functions will not be allowed.

Table 5-5

Data administrator privilege = *CHANGE commands

Command	Description
CHGSCAUT	Change product authority
CHGSRVRDFT	Change server default attributes
CHGSRVRSTS	Change server status
CHKSRVCAT	Check server catalog information
DSPSCAUT	Display product authority
RMVSCAUT	Remove product authority
SETLIBACC	Set library access

Table 5-6 shows the STRATEGY commands, but they should not be used. They're provided for server use and are listed here as reference information only. Using these commands may cause unpredictable results.

Table 5-6
STRATEGY commands

Command
AOSMAST
CACHEMGR
RSCGWY
SCJOBMGR
SCLLQRY
SPBCHPRC
TCPMAST

User Exit Programs

A user exit program can be specified on the iSeries to accept or reject a request before the function is actually carried out. This allows for additional user-defined security not provided by the iSeries or Warehouse Manager. Specify the user exit program name on the client support request access (PCSACC) parameter of the iSeries change network attributes (CHGNETA) command. See the example below.

The PCSACC exit program on the CHGNETA command is supported. *REGFAC is also supported, but if the PCSACC value is *REGFAC the user will need to run the ADDEXITPGM command—with SC_QUERY_ROW_SEC for the Exit point parameter and SCRS0100 for the Exit point format parameter—to use User Exit Programs.

Example:

```
CHGNETA PCSACC(mylib/myuepgm)
```

This tells Client support to call this program before it executes a file transfer, virtual printer, or message function. Warehouse Manager uses the same network attribute and program parameters as Client Access/400 before it prepares each SQL statement or before submitting a batch query.

Following are descriptions of the two parameters that are passed to the user exit program when using Warehouse Manager or Client Access/400. More detailed documentation on user exit programs can be found in the IBM *OS/400 Server Concepts and Administration* document.

- Parameter 1: One byte exit program return code. The request is allowed if the user exit program sets this value to hex F1 (the character "1"). The request is rejected if this parameter is set to any other value.
- Parameter 2: This parameter contains information about the request. The following table describes the structure used by STRATEGY and is identical to the Client Access/400 File Transfer structure.

Table 5-7
Structure of the second parameter passed to a user exit program¹

Value	Type	Description
User profile	Character 10	Contains the user profile making the request.
Application name	Character 10	Contains the name of the PC application making the request. ²
Requested function	Character 10	Contains the type of SQL statement, SELECT, JOIN, INSERT, UPDATE, DELETE, CREATE, DROP, GRANT, etc. ³
Object name	Character 10	Contains the object (table) name. ^{4,5}
Library name	Character 10	Contains the library name or *LIBL. ^{4,5}
Member name	Character 10	Contains the member name, *FIRST or *LAST. ^{4,5} *ALL is not supported by iSeries SQL.
Record format name	Character 10	Contains the record format name. STRATEGY passes a value of *ONLY for the record format name. ⁵
Length of next field	Zoned 5,0	Length of the statement in the following request.
Entire request	char(32500)	SQL statement

1. When developing a user exit program, be aware that Client Access/400 Virtual Print may continue using the first user exit program it calls even though you use CHGNETA to change the user exit program.
2. An asterisk for the first character of the name signifies a STRATEGY Client application. Otherwise, the name is the ODBC Client application name specified by the PC on the ODBC connection string.
3. For SELECT statements involving multiple files, the requested function value is passed as JOIN and the user exit program is called once for each table involved in the join select.
4. The library name does not appear in the SQL statement when *LIBL is implied. The member name never appears in the SQL statement (members are handled through file overrides).
5. Blanks are passed for this field for any SQL statement other than SELECT, INSERT, UPDATE or DELETE.

Distributed Relational Database Architecture (DRDA)

DRDA is the IBM DB2 implementation for database operations on remote DB2s. The correct operation of Warehouse Builder requires that connections to remote iSeries databases be set up correctly. If there are errors with the configuration or if the IBM DDM server is not running, Warehouse Builder sets will not extract and/or write data from one iSeries database to another.

To make sure that remote databases are configured properly:

- ▶ Enter the command WRKRDBDIRE and press Enter.
 There should be a list of remote database names corresponding to other iSeries systems. If the required systems are present, the following steps may be skipped.
- ▶ The first step in adding a remote database entry is to decide which communications protocol to use. You have the choice of SNA or TCP/IP.
- ▶ When configuring the entry for SNA, the APPC information for the remote system is required. Consult the IBM APPC manual to determine and configure the name, the device, the local location, the remote network identifier, and the mode. When configuring the entry for TCP/IP, the TCP/IP address is required. Determine the TCP/IP address of the remote iSeries.

- ▶ Use the command ADDRDBDIRE and fill in the required parameters depending on the communications protocol chosen. Usually, the relational database being added is the system name of the remote iSeries.

If the communications protocol for the remote database entry is TCP/IP, the DRDA server must be started on the remote iSeries system. Use the command STRTCPSVR SERVER(*DDM) on the remote iSeries to ensure that the server is running.

If the communications protocol for the remote database entry is SNA, ensure that the device being used is in the varied-on state and is active. Consult the IBM APPC documentation for instructions, if necessary.

Analyzer OLAP Maintenance

This section discusses typical setup and maintenance procedures that system administrators must perform on the Analyzer OLAP functions of the Warehouse Manager Server. It also includes information about customizing the Analyzer for the Web applets, setting up the iSeries HTTP Server, and other miscellaneous notes.

Configuring the iSeries HTTP Server

This section gives an overview of configuring the iSeries HTTP Server for use with Analyzer for the Web. It is assumed you have a basic understanding of the iSeries HTTP Server. If you need more detailed information regarding the server or the configuration of the server, see *IBM TCP/IP Configuration and Reference Version 3*.

Configuring a New HTTP Server for Analyzer for the Web

If the iSeries HTTP Server on the iSeries has never been configured, at a minimum you must use the following procedure to add basic configuration settings. IBM ships a default configuration file with the iSeries HTTP Server. However, all entries in the file are commented out using the “#” symbol. Figure 5-2 on p. 77 shows a sample HTTP configuration file with the required settings in bold text.

To configure a new HTTP Server:

- ▶ At an iSeries command line, enter CFGTCPHTTP and press Enter.
- ▶ Choose option 2, Work with HTTP Configuration.
- ▶ Specify the configuration name. CONFIG is the default name.
- ▶ Press Enter.

- ▶ Choose Add (option 1) and enter the following above any map entries:

Pass */serverlib/* /serverlib/ProdData/Analyzer/**

where *serverlib* is the name of the library in which Warehouse Manager is installed.

This entry specifies that the HTTP Server is allowed to serve up documents residing in the *serverlib/ProdData/Analyzer* directory. With a Web browser, the end user would need to specify only the *serverlib* as a path to the requested document.

For example, if the Warehouse Manager Server has been installed in a library named *SCSERVER*, enter */SCSERVER/index.html* to bring up the first Analyzer for the Web page. The HTTP Server would actually find the *index.html* page in your system IP or in *alias/SCSERVER/ProdData/Analyzer*.

Note: The URLs may be case sensitive, depending on the configuration. If you want users to be able to enter the URL in lowercase, repeat step 5 and add another entry to the configuration file. The second entry would be the same as the first entry, except that the *serverlib* name would be in lowercase.

- ▶ Press F3 twice to exit the HTTP configuration.
- ▶ Recycle the HTTP Server.
 - To end the server: `ENDTCPSVR SERVER(*HTTP) HTTPSVR(*ALL)`
 - To start the server: `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ALL)`
- ▶ Press Enter.

Note: If you do not restart the HTTP Server, your changes will not take effect. You should now be able to access Analyzer for the Web from your browser by specifying the URL `http://TCPaddress/serverlib/index.html`

where:

- TCPaddress is your iSeries TCP/IP address.
- *serverlib* is the name of the library in which the Warehouse Manager Server is installed.

Note: If, after starting the server, you cannot access or display *index.html*, double-check the HTTP configuration file. Specifically, make sure that the entries that you made are not preceded by the “#” symbol. The “#” symbol is a comment character in the HTTP configuration file and all entries following it are ignored.

The file that contains the HTTP configuration on the iSeries is *QUSRSYS/QATMHTTPC*, member *CONFIG*. You can query against this file to retrieve the settings. A sample HTTP configuration file with the required settings in bold text is shown in Figure 5-2.

Figure 5-2

Sample HTTP configuration file

```

# * * * * *
# iSeries Internet Connection Server Configuration #
# * * * * *
# HostName          your.full.host.name          #
# The default port for HTTP is 80. If you change this use a port number greater than 1024 #
# Port              80                          #
# Enable            GET                          #
# Enable            HEAD                          #
# Enable            POST                          #
# Entries added for STRATEGY Analyzer for the Web #
Pass /showcasecorp/* /scserver/proddata/web/*
Pass /SERVERLIB/* /SERVERLIB/ProdData/Analyzer/*
Pass /serverlib/* /serverlib/ProdData/Analyzer/*
# End of STRATEGY Analyzer for the Web section #
Welcome index.html
# Map      /test/*      /as400/*      #
# Pass     /as400/*     QDLS/400HOME/* #
# Pass     /httpfile/*  QSYS.LIB/AS400LIB.LIB/HTML.FILE/* #
# Pass     /doc/*       /QDLS/graphics/* #
# Pass     /file/*      /www/webdata/* #
# Fail     /QDLS/TESTING/* #
# Redirect /wsg         http://hostname:5061/WSG #
# Exec     /cgi-bin/*   /QSYS.LIB/MYCGI.LIB/* #

# *** DIRECTORY LISTINGS *** #
# DirAccess      On          DirShowSize      On          #
# Welcome        Welcome.html  DirShowBytes     On          #
# Welcome        index.html    DirShowOwner     Off         #
# AlwaysWelcome  On          DirShowDescription On          #
# DirReadme      Top          DirShowMaxDescrLength 25         #
# DirShowDate    On          #

# *** LOGGING *** #
# AccessLog      ACCESSLOG    InputTimeOut     2 mins     #
# ErrorLog       ERRORLOG     OutputTimeOut    20 mins    #
# LogFormat      COMMON      ScriptTimeOut    5 mins     #
# LogTime        LocalTime #
# * * * * *

```

Additional configuration options are available to secure the HTTP Server. For more information on these settings, see *IBM TCP/IP Configuration and Reference Version 3*.

Configuring the Orion Server for Analyzer for the Web

To use Analyzer for the Web, you must install and configure an HTTP server. If you also use Enterprise Reporting, we recommend using Orion Application Server as an HTTP server. It is available during server installation.

To configure Orion Server:

- ▶ Use the iSeries WRKLNK command to open the following XML file for editing:
 - Type WRKLNK OBJ('/scsserver/orion/application-deployments/default/defaultWebApp/orion-web.xml')
 - Select option 2 to open *orion-web.xml*

Note: As a precaution, you may want to create a backup of *orion-web.xml*.

- ▶ Add the following bold line of code to *orion-web.xml* (where *scsserver* is your server library), save the file, and close it.

```
<?xml version="1.0"?>
<!DOCTYPE orion-web-app PUBLIC "-//Evermind//DTD Orion Web Application 2.3//EN" "http://www.orion-server.com/dtds/orion-web.dtd">

<orion-web-app
  deployment-version="1.5.2"
  jsp-cache-directory="/persistence"
  temporary-directory="/temp"
  servlet-webdir="/servlet"
>
  <virtual-directory virtual-path="/showcasecorp" real-path="/SCSERVER/proddata/web/" />
  <virtual-directory virtual-path="/scsserver" real-path="/scsserver/proddata/Analyzer/" />
  <virtual-directory virtual-path="/SCSERVER" real-path="/SCSERVER/proddata/Analyzer/" />
</orion-web-app>
```

- ▶ To load this change, use the iSeries ENDORI command to stop the Orion server and the STRORI command to restart it. For more information about these commands, see Appendix A and the *Enterprise Reporting Administrator's Guide*.
- ▶ To access Analyzer for the Web, use the following URL:
http://iSeries:ORIONPORT/scsserver/index.html

The Orion port number was chosen during Orion installation. For information on installing the Orion Application Server, see the *Enterprise Reporting Administrator's Guide*.

Analyzer for the Web Support for Firewalls and Proxy Servers

Analyzer for the Web requires a direct socket connection to the Warehouse Manager Server to retrieve views and data. In a typical client/server installation, the Analyzer for the Web Java applet communicates with the Warehouse Manager Server using standard TCP/IP over a socket connection. A firewall or proxy server can prevent the applet from establishing a communication link to the server, making it necessary to configure the Analyzer for the Web Java applet to use the standard HTTP protocol in combination with the AOSServlet. The AOSServlet passes the requests from the applet to the Warehouse Manager Server.

Note: You must use the servlet with the IBM WebSphere Server or the Orion Server. The servlet and the server must be running on the same iSeries.

To Configure the AOSServlet with the IBM Websphere Server

- ▶ Add the *ShowCase.jar* file to the Java Classpath for your Web server. The new servlet is packaged as part of the same JAR file.
- ▶ Add the new servlet. Choose a servlet name, such as *AOSServlet*. The class name should be *com.showcasecorp.aosservlet.AOSServlet*. There are no additional servlet properties required.
- ▶ To configure the usage of the servlet interface in your HTML documents, you must change two applet parameters. You will change the *applet.html* document, which can be found in the IFS folder called */serverlib/ProdData/Analyzer/javaClient* where *serverlib* is the name of the library in which the Warehouse Manager Server is installed. Find the following two applet parameters:

```
<PARAM name=UseCGI value="False">
```

```
<PARAM name=CGIURL value="Servlet URL">
```

Change the UseCGI parameter to True and the CGIURL to the URL for your servlet. For example, enter `http://PADDRESS/servlet/AOSServlet` where IPADDRESS is the IP address your iSeries Web Server is running and *AOSServlet* is the same name that you configured in the third step.

To test if the servlet is configured correctly, enter the CGIURL directly into your browser. You should see a message indicating that you have connected to the Analyzer for the Web servlet. If you do not get the message, then your servlet is either not configured correctly or the URL is incorrect.

If the servlet is configured correctly and Analyzer for the Web still is unable to connect to the server, you may have to configure the ServerIP parameter to the applet in your html documents (for example, *applet.html*) as follows:

```
<PARAM name=ServerIP value="localiseriesip">
```

where *localiseriesip* is the native IP address of the iSeries on which Warehouse Manager Server is running.

Note: If you want some clients to use the servlet interface and some to not use the servlet interface, you should have different HTML documents for each type.

To Configure the AOSServlet with the Orion Server

- ▶ Use the iSeries WRKLNK command to open the following xml file for editing:
 - Enter WRKLNK OBJ('/scserver/orion/default-web-app/WEB-INF/web.xml')
 - Select 2 to open *web.xml*

Note: As a precaution, you may want to create a backup of *web.xml*.

- ▶ Add the following bold code to *web.xml* (where TCPaddress is your iSeries TCP/IP address and OrionPort is your Orion port number), save the file, and close it.

```
<?xml version="1.0"?>
<!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.2//EN"
"http://java.sun.com/j2ee/dtds/web-app_2_2.dtd">

<web-app>
  <!-- A demo servlet, add servlets below -->
  <servlet>
    <servlet-name>snoop</servlet-name>
    <servlet-class>SnoopServlet</servlet-class>
  </servlet>
  <servlet>
    <servlet-name>AOSServlet</servlet-name>
    <servlet-class>com.showcasecorp.aoservlet.AOSServlet</servlet-class>
    <init-param>
      <param-name>UseCGI</param-name>
      <param-value>True</param-value>
    </init-param>
    <init-param>
      <param-name>CGIURL</param-name>
      <param-value>http://TCPaddress:OrionPort/servlet/AOSServlet</param-value>
    </init-param>
  </servlet>
</web-app>
```

- ▶ Follow the instructions in the third step of “To Configure the AOSServlet with the IBM Webspere Server,” on p. 79 but when changing the CGIURL to the URL for your servlet, enter *http://TCPaddress:OrionPort/servlet/AOSServlet* where *TCPaddress* is the TCP/IP address your iSeries Web Server is running and *OrionPort* is your Orion port number.
- ▶ Use the iSeries ENDORI command to stop the Orion server and the STRORI command to restart it. For more information about these commands, see “Warehouse Manager Server Commands” on p. 128 in Appendix A and the *Enterprise Reporting Administrator’s Guide*.
- ▶ To test if you’ve successfully configured the AOSServlet, use the following URL:
http://TCPaddress:OrionPort/servlet/AOSServlet
where *TCPaddress* is the TCP/IP address your iSeries Web Server is running and *OrionPort* is your Orion port number.

This URL should display a simple page informing you that the servlet is properly configured and running.

Configuring an Existing HTTP Server for Analyzer for the Web

If the iSeries running the Warehouse Manager Server already has an HTTP Server running, the configuration may be more complex. At a minimum, you must specify to the HTTP Server that the documents in the IFS directory */serverlib* (where *serverlib* is the name of the library in which Warehouse Manager Server is installed) can be served to client browsers.

In some cases, it is necessary to only add the configuration file entries described in the previous section. As long as you insert your pass entry above any map entries, you should not experience any problems. However, in other instances, the HTTP configuration file may contain a map entry pointing to a different IFS directory. In that situation, you must create a symbolic link that makes it appear as though the library containing Analyzer for the Web resides inside the directory specified in the map entry.

Note: If you are using an HTTP Server other than the iSeries HTTP Server, such as I/NET HTTP Server, it may be necessary to create a symbolic link to load the Analyzer for the Web HTML pages.

To Create a Symbolic Link

- ▶ From an iSeries command line, enter the command:

```
ADDLNK OBJ('/serverlib')
NEWLNK('/mapdirectory/serverlib')
LNKTYPE(*SYMBOLIC)
```

where:

serverlib is the library in which Warehouse Manager Server is installed

mapdirectory is the directory specified in the map entry of the HTTP configuration file.

- ▶ Press Enter.

- ▶ Restart the HTTP Server.
 - To stop the server: `ENDTCPSVR SERVER(*HTTP) HTTPSVR(*ALL)`
 - To start the server: `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ALL)`

After creating a symbolic link, Analyzer for the Web users connect by specifying the library name of Warehouse Manager Server (for example, `http://serverlib/index.html`, where `serverlib` is the name of the library in which Warehouse Manager is installed). The HTTP Server automatically adds a prefix for the mapped directory.

Depending on your existing configuration, you may need to make additional changes to the configuration file. See *IBM TCP/IP Configuration and Reference Version 3* for more information.

Using National Language Support

The Warehouse Manager Server must interact with Analyzer using an ASCII CCSID to represent character data. With Essbase/400, DB2/400, and OS/400, the Warehouse Manager Server would use an EBCDIC CCSID. Defaults for both CCSID values are calculated at installation time based on your iSeries system settings. For additional information, including the supported EBCDIC and ASCII CCSIDs, see the “Installing Server Software” chapter in the *STRATEGY Installation Guide*.

The client CCSID setting is used during startup time. It must be a valid CCSID supporting translation to/from the EBCDIC CCSID. This does not affect what code page or encoding windows and web clients are able to run. The Analyzer OLAP part of the Warehouse Manager Server will support multiple windows and web clients on any code page or encoding that supports translation to a specific EBCDIC CCSID. To view and/or change the EBCDIC and ASCII CCSIDs for Analyzer OLAP, see the following sections.

Displaying the Current EBCDIC CCSID

- ▶ To display the current EBCDIC CCSID value, enter the command:

```
DSPFD serverlib/DBDEF
```

where *serverlib* is the name of the Warehouse Manager Server library.

- ▶ Press Enter.
- ▶ Page down to locate the line specifying the coded character set identifier. This is the EBCDIC CCSID.

Changing the EBCDIC CCSID

If you need to change the EBCDIC CCSID value, you must reinstall Warehouse Manager Server and reconfigure the EBCDIC CCSID. See the “Installing Server Software” chapter in the *STRATEGY Installation Guide* for details.

Displaying the Default ASCII CCSID

- ▶ To display the current ASCII CCSID value, enter the command:

DSPPFM *serverlib*/CFG system

where *serverlib* is the name of the Warehouse Manager Server library.
- ▶ Press Enter.
- ▶ Locate the line starting with ClientCCSID=. The value displayed is the default ASCII CCSID.

Changing the ASCII CCSID

- ▶ Add the Warehouse Manager Server library to your library list.
- ▶ From an iSeries command line, enter the command CHGSCSVRA and press Enter.
- ▶ Change the value of the ClientCCSID.
- ▶ Restart the server with the STRSCSVR command. See Chapter 2 for details.

HTML Documents

Several HTML documents are installed when the Warehouse Manager Server is installed. These documents are installed into the IFS directory with the same name as the library in which it is installed.

Note: Analyzer for the Web is a Java-based application. An HTML document must exist before you can use a Java applet.

The Analyzer subdirectory contains sample Web pages. It employs the VMDetector applet to test the capabilities of the user's browser. As a result, the correct Web page locally installs the archives or launches the selected edition. You can use the pages as provided or modify them to meet custom needs.

The following document is installed in `\\serverlib\ProdData\ANALYZER`:

- *index.html* is a home page from which users can launch Analyzer for the Web by linking to pages in the *javaClient* directory.

The following documents are installed in `\\serverlib\ProdData\ANALYZER\javaClient`:

- *launcher.html* inside the *javaClient* directory launches the VMDetector with parameters specifying the splash screen image and version number.
- *applet.html* inside the *javaClient* directory will launch Analyzer for the Web. See "Tailoring Analyzer Web Pages" on p. 85 for more information.
- *default.html* inside the *javaClient* directory acts as the all-purpose launch page for browsers with JDK 1.1 that are not Internet Explorer 4.0 or Netscape Communicator/Navigator 4.05 or higher.

- *troubleShoot.html* notifies the user in the event that the browser is not JDK 1.1-compliant and offers links to the latest versions of Microsoft Internet Explorer or Netscape Communicator/Navigator.

All code for the VMDetector, common images, and archives are stored in the *shared* directory.

Analyzer's directory must be accessible to your Web Server (also known as an HTTP Server). To allow for this, define an additional document directory or a directory link using your Web Server configuration software. For additional information, see "Configuring the iSeries HTTP Server" on p. 75.

Getting Users Started

Zero deployment is a major benefit of the browser client version of Analyzer for the Web. Provide all users with the Analyzer for the Web home page address or a specific launcher page. An example might be:

- //system/serverlib/ProdData/Analyzer/Index
- //system/serverlib/Index.html

Adding Databases

Before users can access databases from Analyzer or Analyzer for the Web, you must specify which databases will be accessible. This is done through Warehouse Manager Client.

To add databases to be accessed with Analyzer or Analyzer for the Web:

- ▶ Start Warehouse Manager Client and connect to the data source containing the Warehouse Manager Server.
- ▶ In the list of objects, double-click Analyzer OLAP.
- ▶ An item called Databases will appear. Right-click on Databases and choose Add Database.
- ▶ In the Essbase System Login dialog box, enter the server IP address, Essbase/400 user ID, and password and click OK. A list of the databases on the server will appear.
- ▶ Select the database(s) that you want to access from Analyzer by pressing Ctrl and clicking the database name.
- ▶ Click OK.

The databases you selected will appear beneath the Databases object in the Warehouse Manager Client window (you may need to expand the Databases object to display the names). These databases are now accessible to users of Analyzer and Analyzer for the Web.

Working with iSeries Profiles

Although by default Analyzer and Analyzer for the Web require users to enter an iSeries profile to connect, it is not necessary to have an iSeries profile to use Analyzer or Analyzer for the Web.

- ▶ From an iSeries command line, run the command CHGSCSVRA.
- ▶ Change the value of the AOValidateUser setting. This setting controls whether users logging on to the server from either Analyzer or Analyzer for the Web must have a valid iSeries profile.
 - If the value is 0, an iSeries log on check is not performed and the log on is passed directly to Essbase/400. A non-zero value causes the server to validate against an iSeries profile. This is the default value.

Note: If you change the ValidateUser value to 0 so that iSeries profiles are not required, you cannot use Warehouse Manager Client to control application security for Analyzer or Analyzer for the Web. Application security requires an iSeries profile. Restart the server with the STRSCSVR command after changing this setting. See Chapter 2 for details.

To simplify your security scheme for Analyzer and Analyzer for the Web, we highly recommend that you use the Integrated Security feature with Essbase/400. For more information, see “Integrated Security” on p. 113 in Chapter 6.

Note: Some Analyzer OLAP features require a valid iSeries user ID and password. These features are not currently used by Analyzer, but they appear in Query and Report Writer. If you want to use scheduled queries or reports or access remote scheduling, users must log on with a valid iSeries profile.

You may turn off integrated security and still have some users enter valid iSeries user IDs and passwords. In that case, only those users will be able to access the Analyzer OLAP functionality in Query and Report Writer.

Specifying an Analyzer Default User Profile

You may find it easier to set up one iSeries profile and use its settings as the default user profile settings. The AODefaultUserProfile parameter in the CHGSCSVRA command specifies an iSeries user profile as the default user profile. The first time a user connects to the server, the default user profile settings are assigned to the new user. For more information about the AODefaultUserProfile parameter, see “CHGSCSVRA Analyzer OLAP-Specific Settings” on p. 130 in Appendix A.

Tailoring Analyzer Web Pages

The sample set of Web pages provided is enough to support most installations. Feel free to copy and modify *index.html* and the pages in the *javaClient* directory with your own style and logos. However, if you decide to modify the pages, rename them so they will not be overwritten during upgrades or patches.

Note: The Web pages, Java classes, and archives in the shared directory should not be modified.

The Java components of specific Analyzer launcher pages—the VMDetector—stored in archives, can be modified by setting parameters passed to them by their associated APPLET tags. The *launcher.html* and *applet.html* files may be customized as long as they retain a valid APPLET tag.

```
<APPLET  
width="0"  
height="0"  
code="VMDetector.class"  
codebase="../shared/VMDetector"  
name="Analyzer VMDetector"  
>  
<PARAM name="Splash" value="../shared/w_inter.gif">  
<PARAM name="BuildNo" value="3.0.3.15902">  
</APPLET>
```

Understanding APPLET Tag Components

An APPLET tag can be broken into two groups of components:

- **Applet modifiers** describe the applet to the browser. The applet modifiers are specified in the HTML standard and are available for use with all applets.
- **Applet parameters** pass values to the applet itself. The applet parameters are specific to the particular applet.

Applet Modifiers

- **WIDTH:** Amount of room the browser will allow the applet to use horizontally.
- **HEIGHT:** Amount of room the browser will allow the applet to use vertically.

Note: In addition to supporting fixed width and height, Internet Explorer and Netscape Communicator can specify the percentage of browser window width and height. For example: WIDTH=99% HEIGHT=80%

- **CODE:** The name of the applet's main class. Changing this value will not load the appropriate classes.
- **CODEBASE:** The URL of the applet's main class.
- **NAME:** Specifies a name for the applet instance. This makes it possible for applets on the same page to find and communicate with each other.
- **ARCHIVE:** The URL of the JAR archive that contains the classes. Netscape browsers download the archive and use the classes within to accelerate startup. When specified in the *applet.html*, it should not be changed.

Note: When classes are installed locally, omit the ARCHIVE modifier.

Parameters

To set parameters for an applet, you specify a name and value. The names and values are specific to the applet and should be documented in its accompanying documentation.

Table 5-8

VMDetector parameters

Splash	The URL of the image to be displayed as the VMDetector is running.
BuildNo	Associates a version number with the VMDetector to be compared to that of the locally installed classes. If you use a custom page to launch the VMDetector, manually update this parameter when you install future versions of Analyzer.

Table 5-9

Analyzer supported parameters

AutoLogin	Default: "False". If set to "True", attempts a log on upon applet start using the values specified by the "UserID", "Password", and "Language" parameters.
UserID	The user name used for auto login.*
Password	The password used for auto login.*
Language	The language used for auto login.*
UseFrame	Default: "True". Causes the applet to run in a resizable frame outside the browser. To turn off this feature, set the parameter value to "False".
HelpItemText	Default: " " The value of this parameter is displayed on the Help menu.
HelpItemURL	Default: " " The URL of the page to load when the user selects the value of HelpItemText from the Help menu.
BGColor	Default: " " The RGB color specification for the background color of the applet. The default value specifies light gray.
ChartLegLimit	Default: " " Specifies how many items will show on the X axis before they are turned to display vertically. For example, if the value is "5" and your chart has six items, all six will be turned so that they display vertically.
Port	The port number to use for connecting to the Warehouse Manager Server. This value must match the configuration of the server.
LoginHint	Default: " " Changes the hint on the Login dialog box to provide a customizable logon message.
TmplImageDirURL	Default: " " The URL of the directory where AOS is storing images. The default specifies the default location.
ServerIP	Default: " " The IP address or the fully qualified DNS name for the server. If the value is specified, the Warehouse Manager Server must reside on that machine.
UseCGI	For firewall and proxy support. See "Analyzer for the Web Support for Firewalls and Proxy Servers" on p. 79.
CGIURL	For firewall and proxy support. See "Analyzer for the Web Support for Firewalls and Proxy Servers" on p. 79.

*Used only if the AutoLogin parameter is set to "True".

About Distributed Archive Files

Advances in Web browser and Java technology make it possible to locally install the classes necessary to run Analyzer for the Web. The advantages of local installation are faster applet startup and less network traffic. The necessary classes will only be reinstalled and redownloaded from the Web server when new versions of Analyzer for the Web are installed. Currently, Analyzer for the Web supports automatic installation of local archive files for use with Netscape Communicator 4.05 and Microsoft Internet Explorer 4.0 or later versions.

Microsoft Internet Explorer and CAB files

Microsoft Internet Explorer uses CAB archives for locally installed classes. Using Microsoft-specific applet parameters, the files are updated by the Microsoft Internet Explorer's Package Manager whenever new versions are located on the server.

In the installed sample, *IESample.html* in the *samples* folder contains the necessary applet tag with the parameters to install the CAB contained in the same folder.

Table 5-10
Microsoft-specific CAB parameters used in *IESample.html*

namespace	Default: "ShowCaseAnalyzer". The name space into which the classes contained in the CAB should be installed. <i>This value should not be changed.</i>
useslibrary	Default: "ShowCaseAnalyzer." The name space of the library classes. <i>This value should not be changed.</i>
useslibrarycodebase	Default: "shared\ShowCase Analyzer.cab". The URL to the Analyzer for the Web CAB file.
useslibraryversion	Default: "3,0,3,15902". The version number corresponding to the version of the CAB files. When installing new versions or patches of Analyzer for the Web, this parameter must be manually updated in custom pages using CAB files. This number will vary.

Uninstall CAB Files

To manually uninstall CAB files through Internet Explorer:

- ▶ Open Internet Explorer, making sure that Analyzer for the Web is not loaded.
- ▶ From the Tools menu, choose Internet Options.
- ▶ In the Internet Options dialog box, select the General tab.
- ▶ In the Temporary Internet Files section, click Settings. The Settings dialog box will appear.
- ▶ To open a folder with all the installed, downloaded CAB files and components, click the View Objects button. The Downloaded Program Files dialog box will appear.
- ▶ Select the Analyzer component that you want to uninstall.

To uninstall the Analyzer classes:

- ▶ In the Downloaded Program Files dialog box, select the component(s) that you want to remove.

- ▶ From the File menu, choose Remove Program File.
- ▶ From the File menu, choose Close.
- ▶ In the Settings dialog box, click OK.
- ▶ In the Internet Options dialog box, click OK.

Netscape Communicator and JAR Files

The sample Web pages use *netscape.html* in the *shared* folder to install new versions of the classes as they are available from your server. The file contains JavaScript that uses Netscape's SmartUpdate to install the JAR files and maintain version information. *As noted earlier, this page should not be changed.* However, *smartupdate.js* (in the *shared* directory) may be used by your custom pages to install the local classes for Netscape browsers. The *NSLauncher.html* sample page demonstrates its use.

Uninstalling JAR Files

Netscape 4.5 users can uninstall the locally installed classes:

- ▶ Open Netscape Communicator, making sure that Analyzer for the Web is not running.
- ▶ From the Edit menu, choose Preferences.
- ▶ From the category list, double-click Advanced.
- ▶ Click SmartUpdate to open the SmartUpdate panel.
- ▶ Select ShowCase Analyzer.
- ▶ Click Uninstall.
- ▶ Click OK.

Sample Pages

The directory contains several pages showing how to customize the pages that launch Analyzer for the Web. Feel free to use and modify these pages to meet your needs.

Table 5-11
Sample pages

<i>IESample.html</i>	Microsoft Internet Explorer only. Illustrates the use of the Microsoft-specific parameters to locally install the CAB and launch the applet.
<i>NoFrame.html</i>	Launches the applet within the browser's window embedded in an HTML page.
<i>NSLauncher.html</i>	Netscape Communicator only. Launches the applet within the browser's window embedded in an HTML page and installs the JAR file as needed.
<i>Analyzer640.html</i>	Launches the applet within the browser's window embedded in an HTML page and sized for a 640 monitor resolution.

Essbase/400 Server Maintenance

This chapter discusses maintenance procedures and operating considerations for Essbase/400 Server. For information about starting and stopping Essbase/400 Server, see Chapter 2.

Understanding the Essbase/400 Configuration

Essbase/400 Server provides two mechanisms for configuring information about the server on the iSeries (AS/400). Users familiar with Essbase servers on other platforms have traditionally used the Essbase configuration file to configure virtually all Essbase Server default characteristics.

Essbase/400 Server uses the Essbase configuration file in combination with Essbase environment settings to completely control the Essbase/400 Server environment. The following sections explain:

- The objects that Essbase/400 installs
- The Essbase/400 environment and how to use the iSeries configuration file
- How the configuration file relates to the Essbase environment settings

Essbase/400 Installation Library and Path

During installation, you can specify an Essbase/400 installation library name. By default, the name is *ESSBASE*. The Essbase/400 installation program uses the installation library name to create both a *QSYS.LIB* library and a root file system directory. Thus, the default root file system directory name is also *ESSBASE*.

For example, if you specified *MYBASE* as the installation library name, the installation program creates a library named *MYBASE* in the *QSYS.LIB* file system and a directory named *MYBASE* within the root file system.

To differentiate between the areas that Essbase/400 uses to store information, we will refer to the library as the *Essbase library* and the root file system directory as the *Essbase directory*.

Essbase/400 QSYS.LIB Objects Created During Installation

Essbase/400 uses the root file system for virtually all user data associated with Essbase application databases. The data in the Essbase library consist of configuration and program information.

The following sections explain which objects are stored in the *QSYS.LIB* file system. If you are familiar with an Essbase file not mentioned here, it can probably be found in the Essbase directory structure within the root file system.

Files

The configuration file is a source physical file named *CFG* and contains a member called *ESSBASE*. In most of the Essbase documentation, the file is referred to as *ESSBASE.CFG*.

Server and application logs are written in ASCII format to IFS. To view the logs from an iSeries command line, use the *DSPESSLOG* command. *DSPESSLOG* can also copy the logs to the *QSYS.LIB* file system.

Miscellaneous Objects

Essbase/400 also includes miscellaneous objects that implement functions specific to the iSeries platform. The miscellaneous objects include user spaces to store environment information, commands to provide user interfaces, a message queue for additional logging support, and a message file that logs iSeries-specific error messages.

Using Essbase/400 Configuration File Settings

The Essbase/400 configuration file contains keywords and values that define characteristics about Essbase/400 Server. The *Essbase Technical Reference*, accessible by choosing Documentation from the STRATEGY menu, provides specific information about each keyword and value that can be used within the configuration file. The Essbase/400 configuration file can contain the same keywords and values used by Essbase on other platforms. Keywords and values not supported by Essbase/400 Server are ignored.

The Essbase/400 configuration file is kept in the iSeries *QSYS.LIB* file system within the Essbase/400 installation library. The name of the configuration file is *CFG*. The settings that are used when the server is started and while it is active are the settings specified within the *ESSBASE* member.

Most Essbase documentation refers to the configuration file as *ESSBASE.CFG*. Where *ESSBASE.CFG* is referenced throughout the documentation, substitute the Essbase/400 equivalent, *essbaselib/CFG* member *ESSBASE* (where *essbaselib* is the Essbase/400 installation library).

When Essbase/400 is installed, the configuration file is created as an empty source physical file unless it already exists. If the file exists, it is converted to a source physical file, maintaining the settings it contains.

If the file is deleted, use the following command to re-create the file:

```
CRTSRCPF FILE(essbaselib/CFG) RCDLEN(240) MBR(ESSBASE)
ACCPHSIZ(*MAX4GB) SIZE(*NOMAX) CCSID(37)
```

where *essbaselib* is the Essbase/400 installation library.

For more information about creating iSeries physical or source physical files, see “Using Configuration Files from a 5250 Display Session” below. For more information about viewing or editing these files, see “Editing or Displaying the ESSBASE.CFG File” below and “Editing or Displaying the ESSBASE.LOG File” below.

Using Configuration Files from a 5250 Display Session

Essbase/400 uses files in the *QSYS.LIB* file system to store information that an administrator or user may want to view or edit using a 5250 display session. The following sections provide iSeries information for displaying or editing the configuration files from a 5250 display session.

Editing or Displaying the ESSBASE.CFG File

To display or edit the *ESSBASE.CFG* file from a 5250 display session, use the iSeries EDTF command. Use the following syntax:

```
EDTF FILE(essbaselib/CFG)MBR(ESSBASE)
```

where *essbaselib* is the Essbase installation library

Editing or Displaying the ESSBASE.LOG File

To display or edit the *ESSBASE.LOG* file from a 5250 display session, use the iSeries EDTF command. To display the contents of the *ESSBASE.CFG* file, use the following syntax:

```
EDTF STMF('/essbaselib/ESSBASE.LOG')
```

where *essbaselib* is the Essbase installation library.

Using Essbase/400 Environment Settings

Each Essbase/400 installation library has unique settings. Use parameters to tailor the Essbase/400 environment. Some parameters are configured during installation. Others are configured only when you need them. The environment parameters are implemented with OS/400 environment variables. Each Essbase/400 installation has its own set of environment variables.

Two Essbase commands allow you to work with Essbase/400 environment variables:

- CHGESSENV allows you to change the standard Essbase/400 environment variables and additional generic Essbase environment variables.
- DSPESSENV displays the currently configured Essbase environment.

See the following sections for details about configuring, displaying, and removing Essbase/400 environment parameters.

Understanding When the Environment Is Used

The Essbase environment is activated when Essbase/400, MaxL/400, or ESSCMD/400 is started. It is important to remember that even though the Essbase environment can be changed while Essbase operations are active, the changes are not active until all Essbase operations are ended and restarted.

Configuring Essbase/400 Environment Parameters

A standard Essbase environment is configured at installation time. After installation, it can be changed with the CHGESSENV command. The following table lists the standard environment parameters and their associated command parameters.

Table 6-1
Standard Essbase environment parameters

Command parameter	Environment variable
LIB	ARBORLIB
PATH	ARBORPATH
SERIALNUM	ESSBASE_SERIAL
IPADDR	ESSBASE_IP_ADDR
THREADS	SERVER_THREADS
EBCCCSID	ESSBASE_EBCDIC_CCSID
ASCCSID	ESSBASE_ASCII_CCSID
SCRSYNTAX	N/A
ALLOBJ	Run with all object authority
CLASSPATH	CLASSPATH

The generic variables in Table 6-2 can also configure Essbase/400 operations.

Table 6-2
Generic Essbase environment variables

Environment variable	Description
ESSBASE_UPPERCASE	When this variable is defined with a non-NULL value, all text output that Essbase generates in messages and log files is displayed in uppercase English. This environment variable is most effective using Essbase on systems outside the United States that do not support lowercase English characters.
LINES	Specifies the number of lines to be displayed by the Help command in the ESSCMD interface.
ESSBASE_JOURNAL	Specifies the library and name of a journal where Essbase will log debug information. Developers may use this information to investigate problems. <i>Note:</i> Use this variable only at the direction of technical support.

Displaying Essbase/400 Environment Parameters

To display the currently configured environment, run the DSPESEENV command. This command displays all of the defined environment variables and their values.

Removing Essbase/400 Environment Parameters

The standard Essbase environment parameters cannot be removed. Once set, they can be changed only to another value. The generic variables can be removed by specifying *REMOVE as the value on the CHGESSENV command.

Understanding the Configuration and Environment Relationship

Because Essbase/400 configures operating characteristics in two ways, it is important to understand how the characteristics relate to one another.

- The Essbase configuration file allows characteristics common among all Essbase platforms to be defined.
- The Essbase environment settings define either characteristics that are unique to Essbase/400, characteristics that have different value ranges, or defaults for Essbase/400.

Server threads can be configured in both places. If a value is configured in both places, the environment setting takes precedence over the configuration file setting. To reduce confusion, it is recommended that server threads be configured using only the environment settings.

Configuring Essbase/400 for Java

Essbase/400 supports Java-related functionality. See the *STRATEGY Installation Guide* for any Essbase/400 requirements to work with Java. The JVMModuleLocation parameter, which is set in the *ESSBASE.CFG* file, is ignored in Essbase/400. Instead, a standard Essbase/400 environment variable called CLASSPATH has been created. CLASSPATH specifies the location of the Java files needed to run with Essbase/400. Essbase/400 will always seek the CLASSPATH default value. Do not change the default value. Any changes must be appended or Java will not work.

The CLASSPATH default value is:

```
\arborpath\java\essbase.jar:  
\arborpath\java\jaxp\parser.jar:  
\arborpath\java\jaxp\jaxp.jar:
```

Using Essbase/400 National Language Support

Essbase/400 interacts with PC client workstations using an ASCII CCSID to represent character data. To represent character data with DB2/400 and OS/400, Essbase/400 uses EBCDIC CCSID. Defaults for both CCSID values are calculated at installation time based on your iSeries system settings. For additional information, including the supported EBCDIC and ASCII CCSIDs, see “Displaying the Current EBCDIC CCSID” on p. 96 and “Displaying the Current ASCII CCSID” on p. 97. To view and/or change the EBCDIC and ASCII CCSIDs for Essbase/400 Server, see the following sections.

Migrating to the Euro CCSID

If you need support for the Euro CCSID, you will need to change the EBCDIC CCSID. See the following table for easy conversions. If the conversion you need is not in the table, see “Changing the ASCII and EBCDIC CCSIDs” on p. 97.

Table 6-3
Euro CCSIDs

Country, region or language group	EBCDIC CCSID	Euro CCSID
United States/Canada/Brazil	37	1140
Germany/Austria	273	1141
Denmark/Norway	277	1142
Finland/Sweden	278	1143
Italy	280	1144
Spain/Latin America	284	1145
United Kingdom	285	1146
France	297	1147
Multinational #5	500	1148
Iceland	871	1149

To make the CCSID change:

- ▶ Type the iSeries command:

```
essbaselib/CHGESSENV EBCCSID(euro CCSID)
```

where *essbaselib* is the name of the Essbase/400 installation library and *euro CCSID* is the Euro CCSID.

- ▶ Press Enter.
- ▶ Stop and restart the server by following the instructions in Chapter 2.

Displaying the Current EBCDIC CCSID

- ▶ To display the current EBCDIC CCSID value, type the iSeries command:

```
essbaselib/DSPESENV INFO (*ALL)
```

where *essbaselib* is the name of the Essbase/400 installation library.

- ▶ Press Enter.
- ▶ Locate the ESSBASE_EBCDIC_CCSID variable. The value displayed is the current EBCDIC CCSID.
- ▶ To display second-level Help text, set the cursor on the Essbase/400 Server Attributes message and press F1.

Displaying the Current ASCII CCSID

- ▶ To display the current ASCII CCSID value, type the iSeries command:

```
essbaselib/DSPESSENV INFO (*ALL)
```

where *essbaselib* is the name of the Essbase/400 installation library.
- ▶ Press Enter.
- ▶ Locate the ESSBASE_ASCII_CC SID variable. The value displayed is the current ASCII CCSID.
- ▶ To display second-level Help text, set the cursor on the Essbase/400 Server Attributes message and press F1.

Changing the ASCII and EBCDIC CCSIDs

When changing Essbase/400 Server CCSID values, consider the potential impact on your Essbase/400 environment. You may have to re-create part of your Essbase/400 environment because certain characters may be interpreted differently once the change is complete.

In most environments, an EBCDIC CCSID change will not impact the interpretation of character data stored within Essbase/400 Server. If the ASCII CCSID is changed and only invariant (universal codes among CCSIDs) alphabetic characters are used within Essbase/400, the change will have little or no impact. However, if the ASCII CCSID is changed and language-specific characters are used within Essbase/400, you will be required to re-create some information.

Before attempting a change, examine the differences between the CCSIDs that you are considering changing and the importance of the differences. Consult the IBM manual *AS/400 National Language Support* to assist in your research. If you still believe a CCSID change is necessary, follow the steps below:

- ▶ Save outlines, rules files, calculation scripts, and report scripts to the client PC.
- ▶ To make a change, type the iSeries command:

```
essbaselib/CHGESSENV ASCCSID(ccsid1) EBCCSID(ccsid2)
```

where:
 - *essbaselib* is the name of the Essbase/400 installation library.
 - *ccsid1* is the new ASCII CCSID.
 - *ccsid2* is the new EBCDIC CCSID.
- ▶ Verify that the change is correct.
- ▶ Reload all outlines, rules files, calculation scripts, and report scripts from the client PC to Essbase/400 Server.
- ▶ Import any Essbase/400 data that was interpreted differently as a result of the CCSID change.

Administering Essbase/400

iSeries users can administer Essbase/400 Server from a command line using MaxL/400.

Traditionally, ESSCMD/400 has provided functionality that is consistent with the PC ESSCMD (ESSCMD and ESSCMDW). In this release of STRATEGY, Essbase/400 supports the MaxL data definition language. MaxL is positioned as the next generation of ESSCMD.

While ESSCMD/400 and the ESSCMD API continue to be supported, many new Essbase features are accessible only through the MaxL/400 interface. Thus, we recommend using MaxL/400 to automate and maintain your Essbase/400 server. This documentation explains how to use MaxL/400 with Essbase/400 and how to use MaxL/400 with existing ESSCMD/400 and ESSCMD API automation. Since Essbase/400 fully supports the MaxL language, refer to the *Hyperion MaxL User's Guide* for more information.

MaxL/400 and ESSCMD/400 use the iSeries standard input, output, and error devices to retrieve input, write output, and write errors, respectively. MaxL/400 uses IFS files exclusively for its input/output files. ESSCMD/400 uses *QSYS.LIB* physical or source physical files for input/output files. Using parameters from the respective commands, these files can be switched from using the device interface to using a file on the iSeries. Thus, an administrator can script and run a sequence of Essbase commands. From this point forward, we will refer to running MaxL/400 or ESSCMD/400 commands with an input file as the *scripted interface* and running commands without an input file as the *nonscripted interface* or *command line*.

An alternative to ESSCMD/400 is the ESSCMD API Toolkit, which provides OS/400 command interfaces and OS/400 error handling for all ESSCMD/400 operations. For more information, see “Using the ESSCMD API Toolkit” on p. 103.

Command Interfaces

You can run commands in script mode, interactive mode, or command mode.

- **Script mode.** Use a scripting language, such as CL, along with commands to automate processing. No shell is invoked, and commands are processed from an input file.
- **Interactive mode.** Enter statements and receive responses on the MaxL/400 or ESSCMD/400 shell or command line. This mode is effective for day-to-day administrative tasks, error/syntax checking, determining status, or checking information on the fly. In ESSCMD/400, interactive and command modes are very similar.
- **Command mode.** Pass MaxL or ESSCMD statements to the respective interpreter and eventually to the Essbase/400. Use the respective command with parameters to pass the statement and use iSeries commands to specify the statement, the output file, and the escape message level.

Using an Input Script File

To use the script interface, place the commands in an input file and specify an output or error file, depending on the script's intent. If you are running a short script or testing a script, you can run the script interactively while directing the output and error information to the workstation. Thus, the script is properly coded and displays what the script did or did not do. While running scripts in this manner is useful for testing, script information is most often directed to a file. This allows you to start scripts, let them run unattended, and easily verify at a later time that they finished correctly.

When an output or error file is specified, the output or error information that is normally displayed on the screen is written to the specified file.

A script file may contain ESSCMD/400 and MaxL/400 commands. The ESSCMD/400 and/or MaxL/400 session must be initiated with the LOGIN command.

MaxL/400 Scripts (ESSMSH Command)

All filenames referenced in the ESSMSH parameters must be in IFS format. The input and output files are IFS files and in ASCII format. Create the input files in the IFS before specifying them in the command.

The syntax for running a script file with the MaxL/400 shell interpreter specifies an input, output, and error file:

```
essbaselib/ESSMSH INFILE(infile)OUTFILE(outfile) ERRFILE(errfile)
```

where:

- *essbaselib* is the library in which Essbase/400 is installed.
- *infile* is the input script file that contains the input script file member.
- *outfile* is the file that contains the output file member.
- *errfile* is the file that contains the error file member.

Following is an example script file containing MaxL/400 commands:

```
login user1 password on 172.20.1.54;
display application all;
display user all;
drop user mea;
create user mea identified by password member of group '!supervisors' comment 'delete me';
display user all;
msh '/essm40ga/drop.xml' larry;
msh '/essm40ga/register.xml' larry;
export database larry.datload all data to data_file '/essm40ga/exp.txt';
display partition all;
display disk volume all;
display macro larry;
display function larry;
display system;
display filter row all;
logout;
exit;
```

ESSCMD/400 Scripts (STRESSCMD)

To create ESSCMD/400 script files, create an iSeries *QSYS.LIB* file and member to contain the Essbase commands. The file can be created as a physical or source physical file. Set the record size between 80 and 240 characters to allow easy iSeries-based editing and viewing. Once the file and member are created, update the file to contain the Essbase commands that ESSCMD/400 will run.

The syntax for running ESSCMD/400 with an input, output, and error file is:

```
essbaselib/STRESSCMD INFILE(inlibrary/infile) INMBR(inmember)
OUTFILE(outlibrary/outfile) OUTMBR(outmember)
ERRFILE(errlibrary/errfile) ERRMBR(errmember)
```

where:

- *essbaselib* is the library in which Essbase/400 is installed.
- *inlibrary* is the library in which the input script file is located.
- *infile* is the input script file that contains the input script file member.
- *inmember* is the input script file member that contains valid Essbase/400 commands.
- *outlibrary* is the library in which the output file is located.
- *outfile* is the file that contains the output file member.
- *outmember* is the output member where ESSCMD/400 screen output is written.
- *errlibrary* is the library in which the error file is located.
- *errfile* is the file that contains the error file member.
- *errmember* is the error file member where ESSCMD/400 error information is written.

Specifying ESSCMD/400 Output and Error Files within ESSCMD/400 Script Files

When creating ESSCMD/400 scripts, several ESSCMD commands let you specify input, output, and error files for a given command. If a command that you are running requires a filename as a parameter, enter the filename in the syntax that Essbase understands.

However, if you specify a file in the *QSYS.LIB* file system as shown in the example, the library and file must exist before running ESSCMD/400. ESSCMD/400 does not create the library and file. ESSCMD/400 will create the specified member if it does not exist. However, if the member does not exist, verify that the MAXMBRS setting on the file will allow the creation of a new member. If the MAXMBRS limit is reached, the member cannot be created.

The files can be created as physical or source physical files. We recommended a record size between 80 and 240 characters for easy iSeries-based editing and viewing. If the ESSCMD/400 output file is not created in advance, the file is created as a source physical file with a data field of 254 characters. The size of the data field makes it difficult to view or edit the output. If the record size is not sufficiently large to store the ESSCMD/400 output record, the output record is truncated. Verify that the file is created with a record size large enough to accommodate the error messages and any failing input records to be logged to the file.

For example, when using the BUILD DIM or IMPORT commands, you can specify a filename to use for recording errors found during the operation. You must enter an IFS path name. You may store the file in the root file system or the *QSYS.LIB* file system.

The following example stores the information in a library named MYLIB, a file named *MYFILE*, and a member named MYMEMBER within the QSYS file system:

```
/qsys.lib/MYLIB.lib/MYFILE.file/MYMEMBER.mbr
```


The following example stores the information in the *DIMBUILD.ERR* file within the Essbase directory for the Sample Basic application database of the root file system:

```
/ESSBASE/app/Sample/Basic/dimbuild.err
```

When a root file system path is specified, the directories must exist before running ESSCMD/400. The file will automatically be created if it does not exist. If the file exists, it is overwritten.

Using the Command Line Interface

To run commands from a command line interface, issue the respective commands from an OS/400 command line. Do not specify any parameters. The respective command interpreter will start.

MaxL/400 (RUNMAXLSTM)

Issue RUNMAXLSTM from an OS/400 command line. F1 command Help explains the possible values for each parameter. Following is the syntax:

```
essbaselib/RUNMAXLSTM COMMAND(command)  
OUTPUT(outfile) ESCLEVEL(escllevel)
```

where:

- *essbaselib* is the library in which Essbase/400 is installed.
- *command* is the MaxL/400 command.
- *outfile* is the file that contains the output file.
- *escllevel* specifies the level of the messages that are returned. F1 Help lists the parameter values.

ESSCMD/400 (STRESSCMD)

Issue the STRESSCMD command from an OS/400 command line. Do not specify any parameters. The ESSCMD/400 interpreter shell starts.

ESSCMD/400 File Differences

ESSCMD/400 provides nearly identical support for the ESSCMD commands for the PC. Although the interface that ESSCMD/400 provides is very similar to the interfaces provided by ESSCMD utilities on other platforms, there are differences in coding an ESSCMD/400 script file.

Creating ESSCMD/400 Output and Error Files

Before creating or viewing a QSYS file that holds error records, determine the EBCDIC server CCSID that Essbase/400 is configured to use:

- ▶ Find the server CCSID. Run the DSPESSENV command to determine the value of the ESSBASE_EBCDIC_CCSID variable. Place the cursor over the Essbase/400 Server Attributes message and press F1.

- ▶ Create the source file in QSYS with the following command:

```
CRTSRCPF FILE(mylib/filename) RCDLEN(length) CCSID(value)
```

where:

- *mylib* is the QSYS library name.
- *filename* is the QSYS filename.
- *length* is the record length. (Specify a length long enough to handle the longest rejected record. Note that tab characters will be expanded to spaces at a ratio of eight spaces per tab character.)
- *value* is the single-byte CCSID that Essbase/400 Server uses. (On servers using double-byte character sets, *value* must be the single-byte CCSID that is associated with the mixed CCSID of Essbase/400 Server.)

Viewing Output and Error Files

If you wrote the error file to the root file system and want to view the error file, use the iSeries command DSPF. For example:

```
DSPF STMF('/path/filename') 
```

where:

- *path* is the location of the file.
- *filename* is the name of the output or error file to view.

If you wrote the error file to the *QSYS.LIB* file system and want to view the error file, use the iSeries command DSPF. For example:

```
DSPF FILE(mylib/filename) MBR ESSBASE
```

where:

- *mylib* is the QSYS library name.
- *filename* is the name of the output or error file to view.

Output Command

For ESSCMD/400, the OUTPUT command directs process information output from the session to an iSeries file.

```
OUTPUT lib/file(member)
```

where *lib/file(member)* is the library and filename of the output file.

For MaxL/400, the OUTPUT command directs process information output from the session to a text file.

```
OUTPUT ('/path/file.ext') 
```

where '*path/file.ext*' is the path and filename of the output text file.

Reloading Rejected Records

You can specify an error file that receives rejected records if a dimension build or data load error occurs. This file can reside in either the root file system or the *QSYS.LIB* file system. If you want to use the error file to reload rejected records, direct it to a stream file in the root file system. If you want to display only the error file, store the file in the QSYS file system. In this situation, tabs in the file are expanded to spaces, making the file easier to read. If you direct the error file to the root file system, you can use the CPYFRMSTMF command to copy the error file to the QSYS file system to read:

- ▶ Create a source physical file using the CRTSRCPF command described previously.
- ▶ Copy the file from the root file system to the QSYS file system using the command:

```
CPYFRMSTMF FROMSTMF('file name') TOMBR('/qsys.lib/mylib.lib/pf.file/member.mbr') TABEXP(*NO)
```

where:

- *file name* is the full path and name of the file in the root file system in which the error records were originally logged (for example, /essbase/app/smstbc/dbcs/dimerr.txt).
- *mylib* is the QSYS library name.
- *pf* is the QSYS filename.
- *member* is the name of the member to create in the QSYS file system (for example, dimerr.mbr).

Getting Help with ESSCMD or MaxL Commands

You can install Hyperion documentation with STRATEGY. For help on MaxL/400 or ESSCMD/400 commands, see the Hyperion *Essbase HTML Documentation Help*. To view this Help on your PC, from the STRATEGY menu, choose Documentation.

Using the ESSCMD API Toolkit

The ESSCMD API Toolkit provides OS/400 command interfaces and OS/400 error handling for all of the ESSCMD/400 operations. The ESSCMD API Toolkit is automatically enabled and sits atop the ESSCMD/400 function, transforming the script-based interface of ESSCMD/400 into an OS/400 command interface.

The *Essbase Technical Reference*, accessible by choosing Documentation from the STRATEGY menu, includes information that allows you to take advantage of the ESSCMD API Toolkit.

For details about individual ESSCMD API commands, see the iSeries online Help. To use the Help, enter an ESSCMD API command, press F4, and then press F1 for help on a parameter.

Overview

Use the ESSCMD API Toolkit to open and maintain an ESSCMD/400 session across multiple OS/400 commands. You can intersperse ESSCMD/400 commands with other OS/400 commands. You can also code your CL program to respond to errors reported by Essbase on a command-by-command basis, using standard OS/400 error- and message-handling mechanisms.

Command Structure

The ESSCMD API Toolkit consists of four base OS/400 commands, along with numerous supplemental commands that are discussed later. OS/400 command level Help is available for each of the base and supplemental commands. The four base commands are:

- LOGINESS
- RUNESSCMD
- LOGOUTESS
- RTVESSSTS

Note: The STRESSCMD command is not an ESSCMD API command. STRESSCMD acts independently from the ESSCMD API commands.

LOGINESS Command

Issue the LOGINESS command before running other ESSCMD API commands. After issuing the LOGINESS command, RUNESSCMD (or a supplemental command) interacts with Essbase/400 Server under the session opened by the LOGINESS command. Multiple logons with LOGINESS are not allowed.

RUNESSCMD Command

RUNESSCMD lets you run any ESSCMD command directly from the OS/400 command line, or from within a CL or other program. With RUNESSCMD, type the underlying ESSCMD/400 command as follows:

```
RUNESSCMD COMMAND('UNLOADAPP MYAPP')
```

Typically, you will use the LOGINESS command to initiate the session before using RUNESSCMD.

LOGOUTESS Command

LOGOUTESS ends the connection and closes the ESSCMD API session by issuing the ESSCMD EXIT command within the regular ESSCMD interface.

RTVESSSTS Command

From within a CL program, RTVESSSTS can retrieve information about a previous ESSCMD API Toolkit command if the session associated with that command has not been closed.

There are three pieces of information that can be retrieved:

- Essbase message number
- Essbase message severity
- Essbase message text

See the online command Help for more information about the data types of these variables and the possible return values.

Auto-logout Considerations

When you use ESSCMD API commands within another program, consider how Essbase's Auto-logout feature will affect the program. Essbase is installed with a default 60-minute auto-logout time, but this may be reduced. This may be a problem in a program such as the following:

```
LOGIN SVRUSR(userid) SVRPW(password)
ENDESSAPP APPNAME(mycube)
SAVLIB DEV(TAP01) ...
STRESSAPP APPNAME(mycube)
```

The SAVLIB operation may take a long time, particularly if loading or changing a tape is required to complete the operation. If the time between the completion of the ENDESSAPP and start of the STRESSAPP command is longer than the auto-logout setting, the STRESSAPP will fail.

Unfortunately, the auto-logout setting applies to all users, so you can't adjust it for the user ID used in the program. If you prefer not to change the time-out for all Essbase users (it can be set to no auto-logout), run LOGOUTESS and LOGINESS again wherever a problem might occur. For example:

```
LOGIN SVRUSR(userid) SVRPW(password)
ENDESSAPP APPNAME(mycube)
LOGOUTESS
SAVLIB DEV(TAP01) ...
LOGINESS userid password
STRESSAPP APPNAME(mycube)
```

Supplemental Commands

In addition to the commands described previously, the ESSCMD API Toolkit provides alternative commands to RUNESSCMD. There's no need to remember the ESSCMD command syntax. The underlying ESSCMD syntax is mapped to OS/400 command keywords.

Each supplemental command corresponds to an ESSCMD/400 command. For example, the ESSCMD API command STRESSAPP corresponds to the ESSCMD/400 LOADAPP command. These OS/400 commands have the look and feel of standard OS/400 commands, complete with prompted keywords, Help, and so on. For those ESSCMD/400 commands with no corresponding ESSCMD API command, use RUNESSCMD.

With ESSCMD/400, choose an application and database using the SELECT command so that subsequent commands (that do not specify an application or database) can operate upon the selected database. The ESSCMD API supplemental commands require that you specify the application and database on the command itself. The appropriate SELECT command is issued implicitly for you.

Most of the supplemental commands let you utilize the full power of the underlying ESSCMD/400 commands. This list contains all of the Essbase ESSCMD API supplemental commands and indicates which corresponding ESSCMD command is issued. For more information about each ESSCMD command, see the ESSCMD documentation in the *Essbase Technical Reference*, accessible by choosing Documentation from the STRATEGY menu.

Note: The IMPESSSQL and BLESSDIM commands do not permit all forms of their respective ESSCMD commands to be issued. In each case, only the SQL form with a server-based rule file is supported, since this is the most common usage. To handle other cases, use RUNESSCMD.

Table 6-4
ESSCMD API supplemental commands

ESSCMD/400 API Command	ESSCMD/400 Command
ADDESSGRPU	ADDUSER
BLDESSDIM	BUILDDIM
CLCESSDFT	CALCDEFAULT
CPYESSAPP	COPYAPP
CPYESSDB	COPYDB
CRTESSAPP	CREATEAPP
CRTESSDB	CREATEDB
CRTESSGRP	CREATEGROUP
CRTESSUSR	CREATEUSER
CRTESSVAR	CREATEVARIABLE
DISESSLGN	DISABLELOGIN
DLTESSAPP	DELETEAPP
DLTESSDB	DELETEDB
DLTESSGRP	DELETEDGROUP
DLTESSUSR	DELETEUSER
DLTESSVAR	DELETEVARIABLE
ENBESSLGN	ENABLELOGIN
GETESSAPP	GETAPPSTATE
IMPESSSQL	IMPORT
STRESSAPP	LOADAPP
LOGINESS	LOGIN
LOGOUTESS	EXIT
LSTESSAPP	LISTAPP
LSTESSDB	LISTDB
LSTESSGRPS	LISTGROUPS
LSTESSGRPU	LISTGROUPUSERS
LSTESSUSR	LISTUSERS
LSTESSVAR	LISTVARIABLES
CLRESSDB	RESETDB
RNMESSAPP	RENAMEAPP
RNMESSDB	RENAMEDB
RNMESSUSR	RENAMEUSER
RMVESSGRPU	REMOVEUSER
RUNESSCLC	RUNCALC
ENDESSSVR	SHUTDOWNSERVER
ENDESSAPP	UNLOADAPP
UPDESSVAR	UPDATEVARIABLE

Output Processing

RUNESSCMD and other ESSCMD/400 API commands each have an OUTFILE keyword. This keyword follows normal OS/400 OUTFILE conventions. If the file does not exist, it will be created with a record length of 240 characters.

Controlling Error Handling When Programming

The Essbase/400 messaging mechanism defines five message severity levels:

- Informational
- Warning
- Error
- Serious
- Fatal

If Essbase returns only *Informational* and/or *Warning* messages, by default an ESSCMD/400 API command will end with a normal OS/400 completion message. Conversely, if Essbase returns an *Error*, *Serious*, or *Fatal* message, the ESSCMD/400 API command issues an OS/400 Escape message.

To override this behavior, use the optional ESCLEVEL keyword on most ESSCMD/400 API commands. For example, you could use the following sequence of commands to delete an existing application and then re-create it by copying another application:

```
LOGINESS myuserid mypassword
DLTESSAPP myapp ESCLEVEL(*SERIOUS)
CPYESSAPP yourapp myapp ESCLEVEL(*ALL)
LOGOUTESS
```

The DLTESSAPP uses ESCLEVEL(*SERIOUS) so the sometimes expected error message doesn't cause the program to fail:

```
(1051030) ERROR: Application myapp does not exist.
```

Note: In the previous example, the Essbase ERROR level message will behave the same way, resulting in an OS/400 Completion message. Serious and Fatal messages will still result in an OS/400 Escape message.

In the previous ESCLEVEL example, the CPYESSAPP command uses ESCLEVEL(*ALL) because the copy should always work after successfully deleting the application. CPYESSAPP doesn't return any Essbase messages, not even an Informational one. By using *ALL, the program will fail if anything other than what is expected (no messages) occurs. See the online Help for more information about the ESCLEVEL keyword.

In addition to the ESCLEVEL keyword, you can use the RTVESSSTS command within your program to retrieve status information about the most recent Essbase message issued. This lets you programmatically retrieve the level of the message (for example, Informational, Warning), the Essbase message number, and the message text. For example, the following message:

```
(1051030) ERROR: Application myapp does not exist.
```

would return 1051030 for the STATUS (message number), "ERROR" for the MSGLEVEL, and "Application myapp does not exist." for the MESSAGE.

See the online Help for RTVESSSTS for more information about the types and sizes of the keywords for RTVESSSTS.

Debugging

It's possible to see the commands that the ESSCMD/400 API passes to the ESSCMD/400 engine. To log each ESSCMD/400 command, issue the following command in the job ESSCMD/400 is running under. This command will also log the password.

```
ADDENVVAR ENVVAR(ESSBASE_CMDAPI_LOG) VALUE('your_library/your_file')
```

Once the environment variable is set, turn off logging with the next command:

```
CHGENVVAR ENVVAR(ESSBASE_CMDAPI_LOG) VALUE(*NULL)
```

Logging can be turned on and off again with the CHGENVVAR command.

You can use all the standard OS/400 built-in debugging support for your CL program, RPG program, or other program that uses the ESSCMD/400 API.

To more easily debug CL programs that use the ESSCMD API, you may want to create your CL program to log CL Program Commands:

```
CRTCLPGM ... LOG(*YES)
```

In addition, change your job to default to not log CL Program Commands by issuing:

```
CHGJOB LOGCLPGM(*NO)
```

The combination of these changes will result in a job log (via DSPJOBLOG) that shows only the commands issued directly by your CL program and not those issued internally by STRATEGY software.

Maintaining Essbase/400 Log Files

While Essbase/400 is running, entries are added to various log files. These files reside in the IFS root file system in the application directories. Over time, these files can grow very large. If the logs are not properly maintained, they may fill up or consume an unnecessary amount of storage space.

To control the size of the Essbase/400 log file, use the OS400DAYSLOGRETAIN parameter in the Essbase/400 configuration file. For information about how to edit the configuration file, see "Understanding the Essbase/400 Configuration" on p. 91. When starting Essbase/400 Server or an Essbase/400 application, the server uses OS400DAYSLOGRETAIN to determine how many days of log records to keep. Deleted records are not recoverable.

Table 6-5
OS400DAYSLOGRETAIN value

Value	Results
None or Invalid	Retain all records
0	Retain records from current session
1	Retain records logged in last 24 hours
2 and greater	Retain records from specified number of days

You can also manually clear the server and application logs. In Essbase Application Manager, from the Server menu, choose Delete Event Log, or from the Application menu, choose Delete Event Log.

Index Cache Size

The index cache size of an Essbase/400 database depends on the size of the database index page. For every 1 kilobyte (KB) of index page size, you can specify up to 125 megabytes (MB) for the index cache size. The following table shows the maximum sizes allowed for the index cache, depending on the size of the index page.

Table 6-6
Maximum sizes for the index cache

Index Page Size	Maximum Index Cache Size
1KB (1024B)	125MB (131072000B)
2KB (2048B)	250MB (262144000B)
4KB (4096B)	500MB (524288000B)
8KB (8192B)	1000MB (1048576000B)

If the index cache size limit is exceeded, Essbase/400 will reset the cache size to the maximum size allowed for the index page size. For example, if you specify a cache size of 1000MB but your index page size is 2KB, Essbase/400 will reset the cache size to 250MB.

Note: Data file cache is not used. Therefore, a value of 0 is correct.

Loading DB2/400 Data into Essbase/400

Essbase/400 data load and dimension build operations are most efficient by using the SQL interface support. The SQL interface support is shipped as part of the base Essbase/400 package; you do not need to install a separate SQL interface module.

Setting Up Your iSeries to Allow SQL Loads

- ▶ Type WRKRDBDIRE and press Enter.
- ▶ If you do not see an entry for *LOCAL, use option 1 to add a *LOCAL entry.

Loading Data Using the SQL Interface Support

- ▶ Using the Essbase Application Manager, go to a Data Prep Editor window, and from the File menu, choose Open SQL.
- ▶ Select your application and database in the Select Server, Application and Database dialog box and click OK.
- ▶ In the SQL Data Sources selection box of the Define SQL dialog box, the *LOCAL Relational Database Directory entry is highlighted. The Connect information should *not* be filled in.

Note: The *LOCAL entry allows your rules files to be portable between Essbase/400 servers on different systems. You can also select the name of your system from the list, but your rules file must always be used only on that system.

- ▶ Use the Select information to specify the SQL statement to load data into the Data Prep Editor. In the From box, enter the names of the tables in the SQL naming convention, as opposed to the iSeries naming convention. In the SQL naming convention, a period (.) is used between a library and filename. For example:

library.filename

You cannot specify a specific OS/400 database file member name in this interface.

- ▶ If you want to add GROUP BY or ORDER BY clauses, you can do so in the Where box. However, you must first specify at least one WHERE condition. If you do not have a WHERE condition to specify, simply use 1 = 1 as a placeholder WHERE condition in the Where box. For example: 1 = 1 ORDER BY Product.
- ▶ Click OK/Retrieve.

Substitution Variables in Rule File SQL Statements

Essbase substitution variables can be used in calc scripts, report scripts, and spreadsheets. With Essbase/400, you can also use substitution variables within the SQL statement specified in a Data Load or Dimension Build rule file. This allows easy reuse of rule files for multiple data files of the same format or for retargeting loads for test versus production libraries. Modify the value of substitution variables as needed, which modifies the behavior of the scripts and rule files that use the variables.

To use substitution variables in the SQL statement, simply specify the variable prefixed with an ampersand (&). If a variable is used in the Where box and the result of the entire WHERE clause after all variables are substituted is all blanks or an empty string, the entire WHERE clause will be omitted from the SQL statement.

Note: If a variable reference does not have a value assigned to it, no substitution will occur and the original variable name along with the ampersand will be passed in as part of the SQL statement.

Following is an example of the message that will be logged if the variable reference is not assigned a value:

```
Local/Appname/Dbname/Userid/Info(1012042) Substitution variable <xy> doesn't exist.
```

Substitution variables can be managed by choosing Substitution Variables from the Server menu in Essbase/400 Application Manager, or via the ESSCMD commands: CREATEVARIABLE, UPDATEVARIABLE, and DELETEVARIABLE. See the Hyperion Essbase documentation for more information about substitution variables.

Flat File Support

Many non-iSeries systems support character-delimited files, also known as flat files. Flat files often use a delimiter character, such as a tab or space, to separate individual fields within a record.

Essbase/400 can load flat files stored in the IFS application/database directory with a *.txt* extension. ESSCMD/400 can also read from flat files from IFS.

The recommended way to load data is to use a load rule with a SQL statement. To load data from DB2 files, specify the appropriate SQL statement to select the fields that you want to load. You can also load a flat file using the same technique. In this case, the standard OS/400 database file contains a single character field. This field is used to hold character-delimited data.

For example, suppose you have a tab-delimited text file on a PC. You can load data directly into Essbase/400 from this file using ESSCMD on the PC. Alternatively, you can copy this file to the iSeries and load the data using SQL. Following are the steps that you would use:

- ▶ Create a database file on the iSeries using the appropriate record length:

```
CRTPF FILE(MYLIB/FILE) RCDLEN(132)
```

- ▶ Using the FTP PUT command, copy the data from the PC to the iSeries:

```
put c:\pcdata.txt mylib/file.file
```

- ▶ The file MYLIB/FILE now contains the tab-delimited data. Using Application Manager, create or edit the load rule to read this data by using the following select statement:

```
SELECT * FROM MYLIB.FILE
```

- ▶ From the Options menu, choose Data File Attributes and use the File Delimiter tab in the Data File Attributes dialog box to set the appropriate delimiter.

Using Warehouse Builder with Essbase/400

Using Warehouse Builder, you can perform several operations against an Essbase/400 database (for example, clear a database, build dimensions, load a database, calculate a database). When using Warehouse Builder to build dimensions or load data, Essbase/400 uses SQL to read records and load them into the database. Essbase validates each record as it is processed, and Essbase/400 may reject records if it encounters an error.

Dimension Build and Data Load Rejected Records

The member names identified by each record must match the member names specified in the Essbase/400 outline. If the member names do not match (or some other error occurs), Essbase/400 rejects the record. In this situation, if the Essbase/400 build plan defined in Warehouse Builder is set to Continue on Error, Warehouse Builder logs the rejected record in the JOBMSGs table and processing continues.

The records causing an error are logged with a MSG_ID of 'TXT0208' (in the case of a data load) or 'TXT0207' (in the case of a dimension build). The record itself is logged in the MSG_HELP field. The object ID for the Essbase/400 build plan is stored in the OBJECTID field. The sequence number of the Essbase/400 build plan identifies the Essbase/400 operations within the Essbase/400 build plan and is stored in SEQ_NUMBER.

Reloading Rejected Records

After fixing your Essbase/400 outline, it is possible to reprocess only the rejected records (as opposed to reprocessing all of the data). To reprocess the rejected records, use the Essbase Application Manager to create a dimension build or data load rule that points to the JOBMSGs table. Following is an example SQL statement that reprocesses the rejected records of an Essbase/400 build plan whose object ID is 101 and whose sequence number is 1:

```
SELECT MSG_HELP FROM SCSEVER.JOBMSGs
WHERE MSG_ID = 'TXT0208' AND OBJECT_ID = 101 AND SEQ_NUMBER =1
```

Note: To reprocess the rejected records for a dimension build, the MSG_ID value is 'TXT0207'.

To determine the correct object ID and sequence number to use, run the queries (*.DBQ) shipped with Warehouse Builder against the JOBMSGs table. The queries are located in the *REPORTS* folder in which Warehouse Builder is installed.

Naming Essbase/400 Server

When using Warehouse Builder with Essbase/400, the value you specify for Essbase/400 Server in the Build Plan dialog box must be either an IP address in numeric format or a valid (or usable) name on the control server iSeries. To check whether a name is valid, enter the iSeries command CFGTCP from an iSeries session window on the control server iSeries and choose option 10. On the screen that appears, you can verify that the name associated with the TCP/IP address for Essbase/400 Server is the same as the name defined on your PC.

For additional information about using Warehouse Builder with Essbase/400, see the Warehouse Builder online Help.

Security

All Essbase security functions are supported by Essbase/400. However, individual Essbase/400 data files reside in the OS/400 Integrated File System (IFS) and may be accessible via other OS/400 interfaces. You should verify that these files are properly secured with respect to the iSeries defined security mechanisms. In particular, the Essbase security file itself, named *ESSBASE.SEC*, is stored within the IFS file system.

ALLOBJ Authority

With a profile that has *ALLOBJ authority, a user has full iSeries access, including all databases, programs, folders, and files. If you want a user profile to have access to certain iSeries objects but do not want to assign *ALLOBJ authority, you can grant users authority access to certain objects and not to others. For example, a profile at the *USER level can be granted access to library ABC, but not to library XYZ.

Essbase security controls user access. Using a profile that has been assigned the *SECOFR class is the easy way to start Essbase/400 Server. However, to enable a user to run Essbase/400 with anything less than *ALLOBJ authority, follow these steps:

- ▶ Create a user profile with the desired security class to be dedicated to starting Essbase. Make sure that Essbase will always be started with this user profile.
- ▶ Type the command GRTESSAUT to grant the user profile all of the necessary authorities to start Essbase and press Enter.
- ▶ Type CHGESSENV and press Enter.
- ▶ Change the Run with All Object Authority field to *NO.
- ▶ To verify if a user profile is authorized to start Essbase, run the CHKESSAUT command.

Integrated Security

Essbase/400 provides support for Integrated Security. By using OS/400 security controls, Integrated Security lets you control who has access to Essbase/400. As a result, your security scheme for Essbase/400, Analyzer, and Analyzer for the Web is greatly simplified.

Warning: If you are setting up Integrated Security for Essbase/400, create at least one Essbase superuser profile that is not an OS/400 user, and be sure this user profile does not get deleted or demoted. If an Essbase system supervisor ID is the same as any Essbase/400 ID, you may lose your supervisory status.

Once you set up the Integrated Security environment between Essbase/400 and OS/400, OS/400 user profile operations will be applied to Essbase/400 security. Then you can perform most routine Essbase/400 user security changes from OS/400 security interfaces or from Warehouse Manager. You will no longer need to add or remove users or change group memberships via the Essbase/400 security interfaces. However, you will still use Essbase/400 security interfaces to establish which applications and databases a group can access.

If you choose *not* to use the Integrated Security feature, you need not do anything. By default, Essbase/400 runs with its own security system, independent of OS/400 security.

Note: Whether or not Integrated Security is enabled, access to DB2/400 for data loading and dimension building will still require OS/400 authorization information to be entered.

Setting Up Essbase/400 Integrated Security

To set up Essbase/400 Integrated Security, you should be familiar with basic Essbase security and basic OS/400 security, including OS/400's group profile support. This includes experience creating users and groups in Essbase/400 using Application Manager and OS/400 using Warehouse Manager.

For detailed information about basic Essbase/400 and OS/400 security, see the *Essbase Database Administrator's Guide* and the *IBM OS/400 Security Reference* document. For more information about Warehouse Manager, see the Warehouse Manager online Help. You should also be familiar with editing the Essbase/400 configuration files. For information about editing the configuration file, see "Editing or Displaying the ESSBASE.LOG File" on p. 93.

Essbase/400 requires that actual Essbase user IDs be defined for many operations. For example, locks on objects such as outlines, rule files, or calc scripts must be associated with an Essbase/400 user ID. For this reason, the Integrated Security feature of Essbase/400 will automatically create and delete Essbase/400 user IDs in response to OS/400 user profile creations, deletions, and other operations.

Essbase/400 Integrated Security can be configured in one of two ways: you may enable *all* valid OS/400 users to also have access to Essbase/400, or you may enable only a *subset* of OS/400 users to have access to Essbase/400.

To implement Integrated Security, follow these steps:

► **Step 1.** Start Essbase/400 Server.

At this point, you must decide which OS/400 user profiles should have access to Essbase. You can give *all* valid OS/400 user profiles (both existing profiles and those created in the future) access to Essbase, or you can give a *subset* of OS/400 users access to Essbase. Regardless of your choice, you will still be able to give non-OS/400 users access to Essbase as well.

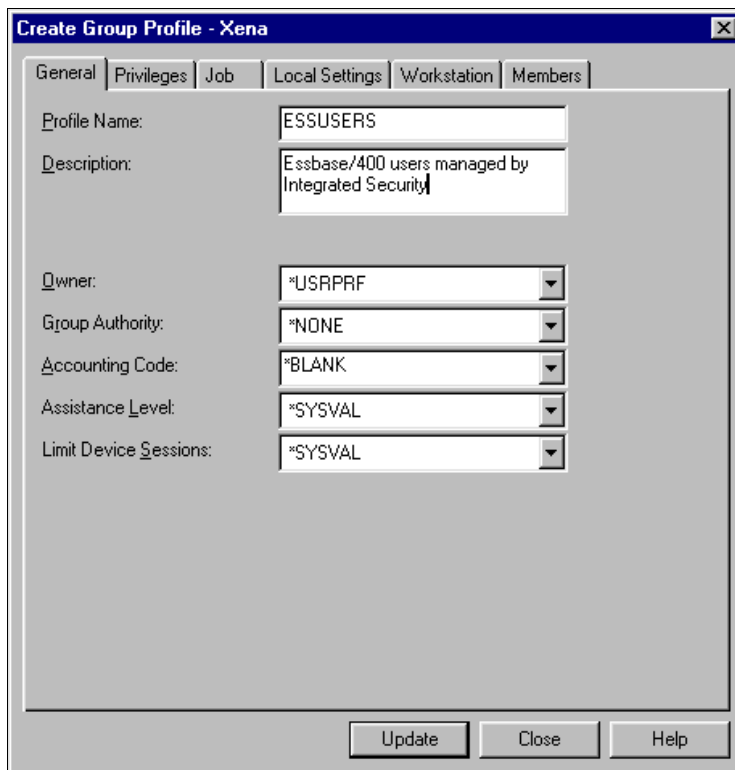
To give *all* current and future OS/400 users access to Essbase/400, skip to step 6. If you want to give only a *subset* of OS/400 users access to Essbase/400, continue with step 2.

► **Step 2.** First, you'll create an OS/400 group, using Warehouse Manager. Open Warehouse Manager. To add a Group User Profile, follow these steps:

- In Explorer, connect to the iSeries data source that you want to create the profile on.
- On the Manage menu, choose Profile, and then choose Create Group. The Create Group Profile dialog box appears.
- Type a name and description for the profile.

The name for the group profile can be anything you want. In this guide, we will assume that you are using the name ESSUSERS.

Figure 6-1
Warehouse Manager Create Group Profile dialog box



- Define any other settings on the other tabs. When setting up Integrated Security, you can accept the defaults that are displayed on these tabs. No special authorities are required.
- Click Update to update the iSeries and refresh the profiles list in Explorer. Then, click Close to close the dialog box.

You have now created a Group User Profile (ESSUSERS) that identifies OS/400 users with access to Essbase/400. Now, you will update the Essbase/400 configuration file to identify the OS/400 group (and corresponding Essbase/400 group, to be created later), which will specify that the OS/400 users have Essbase/400 access.

- ▶ **Step 3.** Without stopping Essbase/400 Server, type the following line in the Essbase configuration file. (For information about editing the configuration file, see “Editing or Displaying the ESSBASE.LOG File” on p. 93.)

```
OS400USERSGROUP ESSUSERS
```

where *ESSUSERS* refers to the name of the OS/400 group you just created.

- ▶ **Step 4.** You can now add OS/400 users to your new group, ESSUSERS. Each OS/400 user profile can be associated with one *primary* Group User Profile and several *supplemental* Group User Profiles. By establishing the ESSUSERS group as either the primary or supplemental group for an OS/400 user, you will indicate that the user has access to Essbase/400.

Note: ESSUSERS can be either the primary *or* supplemental group for any of your users. For example, it could be the primary group for Joe and the supplemental group for Bob.

- ▶ **Step 5.** To associate a Group User Profile with a User Profile, use Warehouse Manager to add users to the ESSUSERS group. How you do this depends on whether ESSUSERS will be the primary or secondary groups for a user.

For those users who *do not* have a primary group:

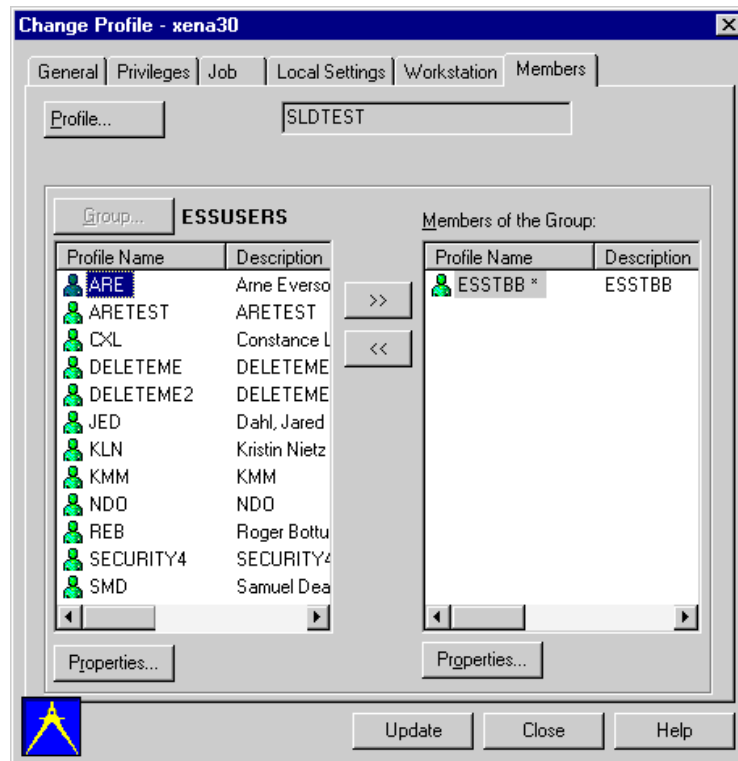
- Find ESSUSERS in the Explorer view. Right-click and choose Change Profile.
- The Create Group Profile dialog box will appear. Click the Members tab and add all the necessary members.
- When finished, click Update.

For those users who already have a primary group, you must add the users individually via the Change Profile dialog box:

- Find the user in the Explorer view. Right-click and choose Change Profile.
- The Change Profile dialog box will appear. Click the Members tab.

Figure 6-2

Warehouse Manager Change Profile dialog box, Members tab



- From the list of all groups, choose ESSUSERS, and then click the double-arrow button to add the group to the user's list of current supplemental groups.
- Click Update. Repeat this for as many users as necessary.

- ▶ **Step 6.** If you have completed steps 2–5, skip to step 7.

You have chosen to allow all valid OS/400 users to access Essbase/400.

To give *all* valid OS/400 users access to Essbase/400, *without* stopping Essbase/400 Server, add the following line to the Essbase configuration file. (For information about editing the configuration file, see “Editing or Displaying the ESSBASE.CFG File” on p. 93.)

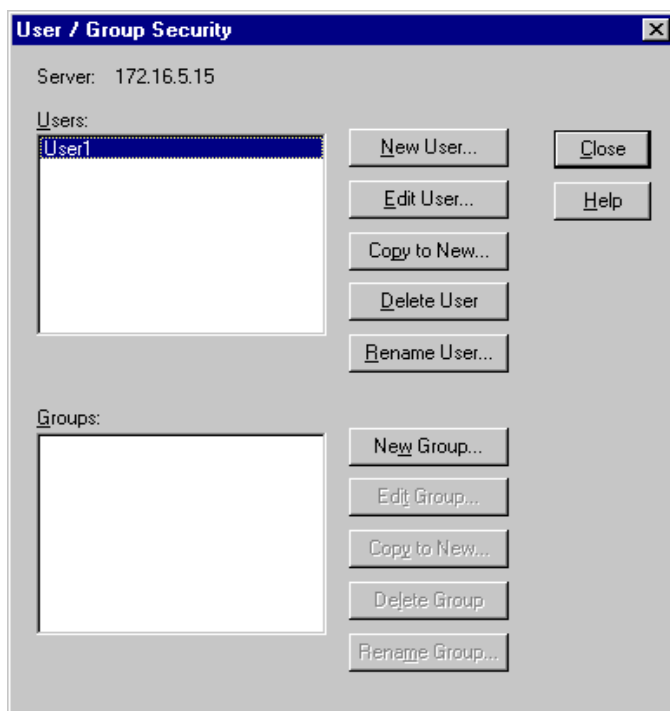
```
OS400UsersGroup *ALL ESSUSERS
```

where ESSUSERS is the group name for Essbase/400.

Note: Be sure to choose a group name that is *not* currently in use by Essbase.

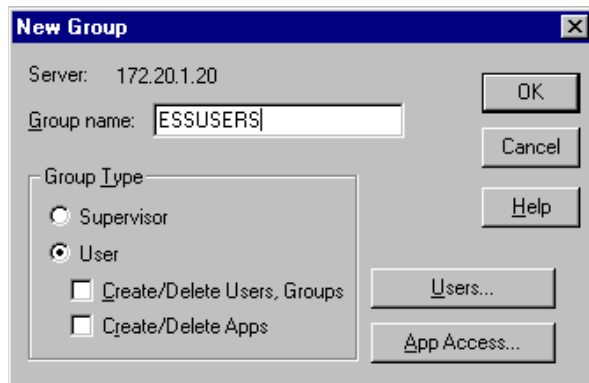
- ▶ **Step 7.** At this point, you have specified the OS/400 users who access Essbase/400 (either a subset of them, or all of them.) Now, using Essbase Application Manager, create the Essbase/400 group (which was specified in the Essbase/400 configuration file in step 3 or step 5) using the following steps:
 - Open Essbase Application Manager on your PC. Connect to Essbase/400 Server. From the Security menu, choose Users/Groups. The User/Group Security dialog box will appear.

Figure 6-3
Essbase User/Group Security



- To create a new Essbase group, click New Group. The New Group dialog box will appear.
- Create an Essbase/400 Group that matches the group name specified in the OS400UsersGroup configuration file entry (in this case, ESSUSERS). Specify the group type as User.

Figure 6-4
Essbase New Group dialog box



- Click OK. The new Essbase group will be added. Click Close to close the User/Group Security dialog box.

Note: Implementing Integrated Security shifts user management to OS/400 and Warehouse Manager. As a result, you will no longer add users to Essbase via Application Manager security. You should use Application Manager security only to create the Essbase group; after that, any user maintenance (for example, adding or deleting users) should be done via OS/400.

- ▶ **Step 8.** Stop and restart Essbase/400 Server. Integrated Security is now in effect.
- ▶ **Step 9.** Using Application Manager, from the Server menu, choose View Event Log. Scroll to the bottom of the log to ensure that there are no errors related to Integrated Security. If you had restarted the server before completing the necessary steps (such as creating the OS/400 or Essbase/400 groups), you may see some error messages. All necessary groups must be created before stopping the server.
- ▶ **Step 10.** When you're finished viewing the log file, from the File menu, choose Close.

Now, as users log on to Essbase/400, Essbase/400 will verify that the user name is an eligible Integrated Security OS/400 user and, if so, the password will be validated against the OS/400 user password. If the OS/400 user is deleted, the Essbase/400 user will be automatically deleted. If the OS/400 user is disabled, the Essbase/400 user will be disabled. Likewise, if the OS/400 password is expired, the Essbase/400 password will be expired. In addition, the OS/400 user profile description will be reflected in the Essbase/400 user description.

At this point, all Integrated Security users will have access only to the Essbase/400 applications and databases to which the Essbase/400 ESSUSERS group has access. For instructions on controlling which OS/400 users have access to which applications and databases, see "Setting Up Group-Level Control" below.

Note: If all Essbase users will have the same level of access, you may choose to modify the rights of the Essbase/400 ESSUSERS group so that as each Integrated Security user is automatically added to Essbase/400, the appropriate rights will be granted via the ESSUSERS group. See the *Essbase Database Administrator's Guide* for information about how to do this. If, however, you want to have different levels of access for different groups, do not give ESSUSERS any additional access; instead, continue with the following section.

Setting Up Group-Level Control

Integrated Security can give you even more control by enabling you to specify at a group level which OS/400 users have access to particular applications and databases. For example, members of the Sales group can see the Sales database, while members of the Human Resources group can see the Human Resources database.

Note: To implement group-level security, you must define an OS400USERSGROUP (such as ESSUSERS). If you haven't done that yet, see "Setting Up Essbase/400 Integrated Security" on p. 114.

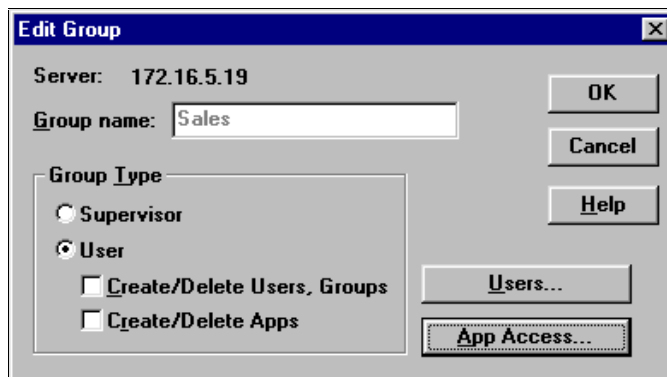
You may already have OS/400 group user profiles established for different departments or roles within your company. For example, you may have an OS/400 group user profile for access to Sales data, one for access to Accounting information, and another for access to HR data. If these group user profiles provide an appropriate method of defining who should have access to Essbase/400 applications, you can use Integrated Security so that as users are added or removed from the OS/400 groups, access to Essbase/400 applications and databases will be adjusted accordingly. If you already have these groups defined, skip the first step. Remember that members of these OS/400 groups must also be members of OS/400 group ESSUSERS (or the equivalent) to participate in Essbase/400 Integrated Security.

- ▶ **Step 1.** If OS/400 groups are not appropriately defined for your Essbase/400 application and database security requirements, you must revise the OS/400 group definitions. Open Warehouse Manager. To add a Group User Profile:
 - In Explorer, connect to the iSeries data source that you want to create the profile on.
 - From the Manage menu, choose Profile, and then choose Create Group. The Create Group Profile dialog box appears.
 - Enter a name and description for the profile.
 - The name for the group profile is your choice. In this guide, we will assume that you are creating two new groups, Sales and HR.
 - Define any other settings on the General through Workstation tabs. These settings will not affect Essbase/400 usage rights. Generally, you can accept the defaults.
 - Click Update to update the iSeries and refresh the profiles list in Explorer. Then click Close.
 - Repeat the first four substeps of this step for each group that you want to create.
 - You can add users to these groups now or later. For users who do not have a primary group, find ESSUSERS in the Explorer view. Right-click and choose Change Profile. The Create Group Profile dialog box will appear. Click the Members tab and add all the necessary members. When you're finished, click Update. Remember that members in these groups must also be members of ESSUSERS or the equivalent.

- ▶ **Step 2.** At this point, you have created Group User Profiles (Sales, HR) that identify users who can access the Sales and HR applications and databases. Now, you will create these groups in Essbase/400:
 - Open Essbase Application Manager on your PC. Connect to Essbase/400 Server. From the Security menu, choose Users/Groups. The User/Group Security dialog box will appear (see Figure 6-3).

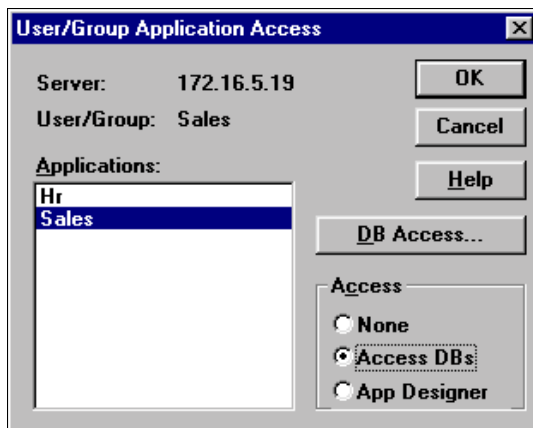
- To create a new Essbase group, click New Group. The New Group dialog box will appear.
 - Create an Essbase/400 Group that matches the name of the group you just created in Warehouse Manager (in this case, HR). Specify the group type as User.
 - Click OK. The new Essbase group will be added.
 - Repeat the previous four substeps to create the Sales group.
- **Step 3.** Adjust the application and database access as needed for each group. For example, the Sales group should have access to the Sales database and the HR group should have access to the HR database.
- In Application Manager, in the User/Group Security dialog box (see Figure 6-3), select a group in the Groups box and click Edit Group. The Edit Group dialog box will appear.

Figure 6-5
Essbase Edit Group dialog box



- In the Group Type box, select User.
- To specify the group's application access, click App Access. The User/Group Application Access dialog box will appear, as shown in Figure 6-6.

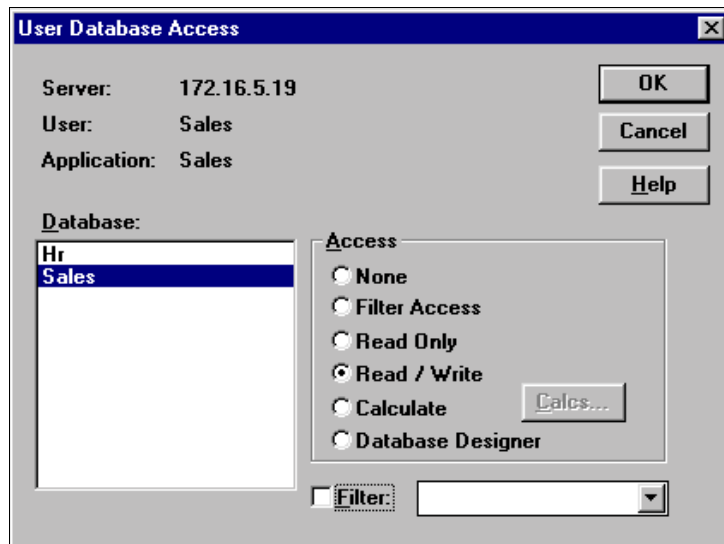
Figure 6-6
Essbase User/Group Application Access dialog box



- Choose an application in the Applications box and select the access in the Access box. If you have created databases under the application, the DB Access button will be enabled if you select

Access DBs. To set the database access, click DB Access. The User Database Access dialog box will appear.

Figure 6-7
Essbase User Database Access dialog box



- Choose a database in the Database box and select the access level for the group in the Access box.
 - When you're finished, click OK. The User/Group Application Access dialog box will appear (see Figure 6-6). Click OK. The Edit Group dialog box will appear. Click OK. The User/Group Security dialog box will appear (see Figure 6-3). Click Close.
 - Repeat the previous six substeps for each group for which you want to set access.
- **Step 4.** Add entries to the Essbase configuration file to tell Essbase/400 which Essbase/400 groups are participating in Integrated Security. Without stopping Essbase/400 Server, add one line, as shown below, to the Essbase configuration file for each group you have defined. (For information about editing the configuration file, see “Editing or Displaying the ESSBASE.LOG File” on p. 93.)
- ```
OS400COMMONGROUP GROUPNAME
```
- where *GROUPNAME* refers to the name of the Essbase/400 group you just created—for example, Sales or HR.
- **Step 5.** Stop and restart Essbase/400 Server.
  - **Step 6.** Using Application Manager, from the Security menu, choose View Event Log. Scroll to the bottom of the log to ensure that there are no errors related to Integrated Security. If you had restarted the server before completing the necessary steps (such as creating the OS/400 or Essbase/400 groups), you may see error messages. All necessary groups must be created before stopping the server.
  - **Step 7.** When you finish viewing the log file, from the File menu, choose Close.

You have completed setup for Integrated Security.

For users who are members of the OS/400USERSGROUP (that is, ESSUSERS), membership in the “common” Essbase/400 groups (for example, Sales, HR) will be dynamically synchronized with the current membership status in the corresponding OS/400 common group.

Automatic creation of an Essbase/400 user occurs when the user first logs on to Essbase/400. Automatic deletion of users (if required) and synchronization of group memberships occurs at every logon, at Essbase/400 startup, and when the Security dialog box in Application Manager is opened.

### ***Troubleshooting Integrated Security***

If you’ve followed all of the previous instructions and Integrated Security still doesn’t seem to be working, check this list for common problems and their fixes. After you fix any problems, you must stop and restart Essbase/400 Server to apply your changes.

- Your log file has any of the following error messages:

1180002: OS400USERSGROUP (ESSUSERS) is not an OS/400 Group User Profile, OS400USERSGROUP disabled

1180003: OS400USERSGROUP (ESSUSERS) is not an Essbase/400 group, OS400USERSGROUP disabled

1180004: OS400COMMONGROUP (HR) ignored. It is not an OS/400 Group User Profile.

1180005: OS400COMMONGROUP (HR) ignored. It is not an Essbase/400 group.

Verify that you have created the group profile in Warehouse Manager, entered the group in the Essbase configuration file, and added the group in Application Manager. Remember that the group name must be the same in all three locations.

- You’ve set up your special groups, HR and Sales, but the users are rejected when they try to log on.

Verify that the users are also members of the OS/400 ESSUSERS group. If you are enabling a subset of your users to access Essbase, those users must first be members of OS/400 ESSUSERS. Then you can add them to any additional OS/400 groups (for example, Marketing, Sales) that you define.

- Your error file includes the following message:

1180008: OS400USERSGROUP (ESSUSERS) has been deleted, OS400USERSGroup disabled

Verify that the OS400USERSGROUP is still defined in Application Manager. If the group has been deleted, it must be re-created for Integrated Security to work properly.

### ***Automatic Deletion of Essbase/400 Users***

There are two ways that Essbase/400 users will be automatically deleted:

- If you have specified only a group name (for example, ESSUSERS) in the OS400UsersGroup configuration entry (that is, you did not specify \*ALL), then any user who is found to be in the Essbase/400 group specified on the configuration entry but is not a member of the OS/400 group, will be automatically deleted from Essbase. This ensures that when an OS/400 user is deleted, the Essbase/400 user will also be deleted.
- If you have specified \*ALL on the OS400UsersGroup entry, automatic deletion will occur when the OS/400 user profile is deleted.

### ***What about Passwords?***

Passwords for Essbase/400 user IDs created by the Integrated Security feature must be changed by changing the OS/400 password via the CHGPWD or CHGUSRPRF command. However, some Essbase/400 Client programs (including Application Manager and Spreadsheet Add-ins) will allow the end user to change a password at Essbase/400 logon time if the password has expired. For an Integrated Security User, this will change the OS/400 password as well.

*Note:* Changing a password via the Essbase/400 Application Manager Security menu option or via ESSCMD will not affect the OS/400 password.

Essbase/400 normally requires that passwords be at least six characters long. OS/400 allows a wide variety of password rules to be optionally enforced and even allows a user exit for password screening. When using the Integrated Security feature, the normal Essbase/400 requirement of six-character passwords is no longer enforced. This will allow existing OS/400 users with shorter passwords to use Essbase without needing a new password. Note that any rules or exits defined under OS/400 will also apply to Integrated Security users when they attempt to change an expired password via an Essbase interface.

### **Changing the Password with a STRATEGY Product**

If the user's Essbase/400 password was changed or expired, the user can update the applicable STRATEGY control table settings with the SCCHKPWD command. The SCCHKPWD command searches the local system and lists and/or updates the necessary control table settings with the new password. For more information about the SCCHKPWD command and its ESSBASE parameter, see "SCCHKPWD Command Parameters" on p. 134 in Appendix A.

### ***Renaming Users and Groups***

When Integrated Security is enabled, certain renaming operations within Essbase/400 are not permitted. The special-purpose groups specified within the Essbase configuration file entries for OS400UsersGroup and OS400CommonGroup may not be renamed. In addition, any Essbase/400 integrated user may not be renamed. These restrictions prevent the Integrated Security configuration from becoming inconsistent.

## ***Essbase/400 Jobs***

You can use the WRKACTJOB command to verify that your Essbase/400 Server is running. The server job will have a name identical to the job description name, a job type of BCH, and a function of PGM-ESSBASE. Applications run under separate jobs with the job name identical to the application name and a job type of BCI.

## SQL Drill-Through Using Essbase Add-ins

When creating a SQL Drill-Through profile, you must qualify an Essbase/400 table with its library. For example, simply entering tablename will fail. Instead you must enter libname.tablename.

For complete instructions about using the SQL Drill-Through feature, from the Windows Start menu, choose

```
Programs
 ShowCase STRATEGY
 Documentation
```

and then click SQL Drill-Through.

## Message Logging

When running Essbase Application Server on platforms other than the iSeries, messages that track the requests and subsequent error messages are displayed on the terminal. Because of internal restrictions on the iSeries, most of these messages are recorded in the Essbase/400 log file and are not displayed on the terminal. To access the Essbase/400 log file on the iSeries, use the DSPESSLOG command.

For information about maintaining the Essbase/400 log file, see, “Maintaining Essbase/400 Log Files” on p. 108.

## Backing Up Essbase/400 Files

As with Essbase running on other platforms, Essbase/400 does not provide direct backup capability. Rather, it assumes that the underlying operating system will provide adequate backup mechanisms.

You can use the SAVLIB or SAV commands to save Essbase/400 applications and libraries. The following procedure creates two save files in QGPL, which hold the QSYS file system and IFS file system information important to Essbase/400. This procedure assumes that Essbase/400 is installed in the default library of ESSBASE. Substitute the name of your installation library for ESSBASE if you did not install Essbase/400 into the default library.

- ▶ Shut down Essbase/400 Server before backing up Essbase/400 files.
- ▶ Type CRTSAVF FILE(QGPL/SAVEQSYS) and press Enter.
- ▶ Type SAVLIB LIB(ESSBASE) DEV(\*SAVF) SAVF(QGPL/SAVEQSYS)  
or  
SAV DEV('/QSYS.LIB/QGPL.LIB/SAVEQSYS.FILE') OBJ('/QSYS.LIB/ESSBASE.LIB')  
and press Enter.
- ▶ Type CRTSAVF FILE(QGPL/SAVEIFS) and press Enter.
- ▶ Type SAV DEV('/QSYS.LIB/QGPL.LIB/SAVEIFS.FILE') OBJ('/ESSBASE') and press Enter.



*Note:* All files found in the *APP* directory, as well as the *ESSBASE.SEC* file, must be backed up and restored together to ensure consistency.

## ***Restoring Essbase/400 Files***

If you used the procedure described in “Backing Up Essbase/400 Files” on p. 124, use the following procedure to restore your Essbase/400 files:

- ▶ Type `RSTLIB SAVLIB(ESSBASE) DEV(*SAVF) SAVF(QGPL/SAVEQSYS)`  
*or*  
`RST DEV('/QSYS.LIB/QGPL.LIB/SAVEQSYS.FILE') OBJ('/QSYS.LIB/ESSBASE.LIB')`  
and press Enter.
- ▶ Type `RST DEV('/QSYS.LIB/QGPL.LIB/SAVEIFS.FILE') OBJ('/ESSBASE')` and press Enter.



# Commands

This appendix lists many of the common iSeries (AS/400) commands, Warehouse Manager Server commands, and Essbase/400 commands you will need for maintaining your iSeries software.

## iSeries Commands

Table A-1 lists common iSeries commands for maintaining a Warehouse Manager library.

Table A-1

*Common iSeries commands*

| Command    | Description                                                | Usage                                                                                                                                                                                                                                                                                     |
|------------|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ADDLIBLE   | Add a library to your library list.                        | To add a Warehouse Manager library to the library list when you need to run a Warehouse Manager command from the iSeries command line.                                                                                                                                                    |
| CFGTCP     | Configure TCP/IP.                                          | To configure iSeries TCP/IP.                                                                                                                                                                                                                                                              |
| DSPPTF     | Display PTFs.                                              | To find the latest PTFs installed on your system.                                                                                                                                                                                                                                         |
| DSPSYSVAL  | Find system information.                                   | To find system information, such as model and serial number. For example, to find the model number, enter DSPSYSVAL SYSVAL(QMODEL); to find the serial number, use QSRLNBR for the system value.                                                                                          |
| DSPUSRPRF  | Display user profile.                                      | To display a user's profile.                                                                                                                                                                                                                                                              |
| EDTF       | Edit file.                                                 | To display or edit iSeries files.                                                                                                                                                                                                                                                         |
| WRKACTJOB  | View active jobs.                                          | To see if the Warehouse Manager Server is running or to view active jobs, such as Essbase/400 and Warehouse Manager jobs.                                                                                                                                                                 |
| WRKOBJLCK  | Check for active jobs against a Warehouse Manager library. | To check for iSeries jobs that may be active against a Warehouse Manager library; in particular, to check for TCP/IP jobs. Useful when installing over the top of an existing Warehouse Manager library or debugging system-related problems (for example, when TCP/IP connections fail). |
| WRKSRVTBLE | View port numbers already in use.                          | To find an available port number when configuring TCP/IP.                                                                                                                                                                                                                                 |

## Warehouse Manager Server Commands

The following tables list the common iSeries Warehouse Manager Server commands for maintaining a Warehouse Manager library. For related security information, see “Command Authorities” on p. 71 in Chapter 5.

*Note:* Before using these commands, add the Warehouse Manager library to your iSeries library list and make sure that you have only one Warehouse Manager library in your library list. Run the command ADDLIBLE serverlib, where serverlib is the name of the library running the Warehouse Manager Server that you want to work with.

Table A-2

Common Warehouse Manager Save/Restore commands

| Command    | Description                 | Usage                                                                                                                                    |
|------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| SAVSRVRINF | Save server information.    | To back up the Warehouse Manager Server or move data stored on the server to another library on the same iSeries or a different iSeries. |
| RSTSRVRINF | Restore server information. | To restore the server data saved using the SAVSRVRINF command.                                                                           |

Table A-3

Common Warehouse Manager TCP/IP commands

| Command   | Description                                                                             | Usage                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHGSCSVRA | Add TCP/IP connectivity support or control Analyzer OLAP functionality (see Table A-7). | To enable or disable TCP/IP support after installation. Requires *ALLOBJ, *SECADM, *IOSYSCFG, *JOBCTL, and *SAVSYS authorities.<br><i>Note:</i> Before using the CHGSCSVRA command, end the server with ENDSCSVR. After using the CHGSCSVRA command, restart the server with STRSCSVR.                                                             |
| ENDORI    | End Orion server.                                                                       | To end the Orion server. Requires *ALLOBJ and *JOBCTL authorities. For more information about the Orion server, see the <i>Enterprise Reporting Administrator's Guide</i> .                                                                                                                                                                        |
| ENDSCSVR  | End Warehouse Manager Server.                                                           | To end TCP/IP. For example, if you need to restart TCP/IP, you should end TCP/IP first. Requires *ALLOBJ and *JOBCTL authorities.                                                                                                                                                                                                                  |
| INSTORI   | Install Orion server.                                                                   | To install the Orion server. For more information about the INSTORI command, including parameters, see the <i>Enterprise Reporting Administrator's Guide</i> .                                                                                                                                                                                     |
| MIGSCREP  | Migrate STRATEGY objects and documents to 4.5 Web deployment framework.                 | To migrate STRATEGY 3.0, 3.5, 4.0, or 4.1 Information Directory objects and documents to the STRATEGY 4.5 Web deployment framework. Requires QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed). For more information, see the <i>Enterprise Reporting Administrator's Guide</i> . |

Table A-3  
Common Warehouse Manager TCP/IP commands (Continued)

| Command   | Description                                                                  | Usage                                                                                                                                                                                  |
|-----------|------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RSTSMPLB  | Restore sample database.                                                     | To install the sample database if not previously installed, or to restore the sample database back to its original structure (if database files were altered or destroyed).            |
| SETJDEOWA | Enable/Disable non-iSeries profile sign-on for J. D. Edwards OneWorld users. | To enable/disable J. D. Edwards OneWorld connections for users who do not have an iSeries user profile. For more information, see "Non-iSeries Profile Sign-On" on p. 52 in Chapter 4. |
| STRORI    | Start Orion server.                                                          | To start the Orion server. Requires *JOBCTL authority. For more information about the Orion server, see the <i>Enterprise Reporting Administrator's Guide</i> .                        |
| STRSCSVR  | Start Warehouse Manager Server.                                              | After an IPL, ENDSCSVR, or after configuring TCP/IP with CHGSCSVRA. Requires at least *JOBCTL authority.                                                                               |
| STRTCP    | Start iSeries TCP.                                                           | To start iSeries TCP/IP. iSeries TCP/IP must be started before you can enable or start TCP/IP support.                                                                                 |

Table A-4  
Common Warehouse Manager Server information commands

| Command    | Description                               | Usage                                                                                                                                                                                        |
|------------|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHGSCAUT   | Change/enter Warehouse Manager passwords. | To enter license information.                                                                                                                                                                |
| CHKSRVRCAT | Check server catalog information.         | To inspect Warehouse Manager catalogs and delete obsolete Warehouse Manager settings.                                                                                                        |
| CHGSRVRSTS | Change server status.                     | To change the status of the Warehouse Manager Server library (for example, SCSEVER) to and from a restricted state. A restricted state prevents users from connecting to the server library. |
| DSPSCAUT   | Display Warehouse Manager authority.      | To see the version of Warehouse Manager installed and the products and licensing information.                                                                                                |
| DSPSCINF   | Display STRATEGY hints and tips.          | To locate hints and tips provided by ShowCase, a division of SPSS Inc.                                                                                                                       |

**Table A-5**  
Common Warehouse Manager Server options commands

| Command    | Description                                                                 | Usage                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHGSRVRDFT | Change server defaults.                                                     | To change the default *PUBLIC authority to all libraries or tables and enforce Warehouse Manager security rules on *ALLOBJ profiles.                                                                                                                                                                                                                                                                                                                            |
| SETLIBACC  | Set library list access.                                                    | To restrict users so they can access only the libraries in their iSeries library lists, or to give users access to all libraries on the iSeries to which they are authorized.                                                                                                                                                                                                                                                                                   |
| SCCHKPWD   | Manage iSeries passwords stored in STRATEGY control tables (see Table A-8). | Update the STRATEGY control tables with a user's new password, or set the UPDATE parameter to *NO to simply generate a list. Type the SCCHKPWD command on the iSeries command line, and press F1 for parameter information.                                                                                                                                                                                                                                     |
| CMPSJAR    | Compile STRATEGY Java files.                                                | The Warehouse Manager Server uses Java. STRATEGY will install several required .jar files in your Integrated File System (IFS). To reduce server start-up time, compile these files on the system running Warehouse Manager Server before you start the server. Compiling the files may affect system performance during the compile. To prevent this, you can compile immediately after the installation finishes or enter a date and time to begin compiling. |

**Table A-6**  
Warehouse Manager Server commands for Warehouse Builder

| Command | Description               | Usage                                                                                                                                                                               |
|---------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STRDD   | Start data distribution.  | Runs a Warehouse Builder distribution set immediately. Can run by set name (OBJNAME) or ID (OBJID). For example, STRDD OBJNAME ('Sales') runs the Sales distribution set.           |
| SBMDD   | Submit data distribution. | Submits a Warehouse Builder distribution set to be run. Can run by set name (OBJNAME) or ID (OBJID). For example, SBMDD OBJID(51) runs the Sales distribution set with an ID of 51. |

## ***CHGSCSVRA Analyzer OLAP-Specific Settings***

Before running CHGSCSVRA, stop the Warehouse Manager Server with ENDSCSVR. Analyzer, Analyzer for the Web, Analyzer Designer, Query, and Report Writer use part of the server as a common repository for storing various objects, including queries, reports, views, forms, pins, pinboards, and lists. Additionally, the server loads the view definitions, retrieves the data from Essbase/400, and packages and sends the data to the Analyzer for the Web Client. Table A-7 contains CHGSCSVRA settings for maintaining this functionality.

Table A-7  
CHGSCSVRA settings for Analyzer functionality

| Setting                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AOEncrypt                    | Enables data stream encryption. The default is encryption enabled. This is a Boolean value with 1 being “enabled” and 0 being “disabled.”                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| AOAllowNewUserDatabaseAccess | Enables a prompt for username and password when the user ID and password used to log on to Analyzer aren’t valid for a particular Essbase/400 database. The default is no prompting. This is a Boolean value with 1 being “enable prompting” and 0 being “disable prompting.”                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| AOSavePasswordMode           | Enables saving of the default preferences whenever users who don’t currently have preferences set log on to the Server. Setting the preferences enables them to be viewed in Warehouse Manager. The default is not to save these settings. This is a Boolean value with 1 being “save the default settings” and 0 being “do not save the settings.”                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| AODefaultUserProfile         | Used the first time a user connects to the Server. The user settings associated with the default user profile, including the system manager setting, are assigned to the new user. You can also (optionally) reset the settings for all existing users.<br><i>Note:</i> If you change the existing default user’s settings and want to apply them to the current users, you must run the CHGSCSVRA command and set the Update All Users option to *YES.<br>Enter a blank (”) profile to turn off the use of the default profile for new users.<br>When specifying default user profile, a conditional prompt asks whether to update the settings for all users. Choosing this option will apply the user settings associated with the default profile to all users. If users are connected to the server at the time this command is entered, their user settings will not be updated.<br>The AOValidateUser parameter cannot equal 0 if you want to use the AODefaultUserProfile setting. |
| AOAllowDataExport            | Disables or enables the data export functions within Analyzer for an individual user or for all users, including e-mail, Lotus Notes, Excel, Text File, Lotus 123W, and the Windows clipboard. If the AOAllowDataExport setting is not specified, the default is to always allow data exporting.<br><i>Note:</i> This setting affects only existing users who have previously signed on to Analyzer. To apply this setting to new users as they sign on, turn on the AODefaultUser option using CHGSCSVRA. The new settings will be returned to the Analyzer Client when the user signs on. It will disable the appropriate commands at that time.                                                                                                                                                                                                                                                                                                                                         |
| AOAllowPrint                 | Disables or enables the print functions within Analyzer for an individual user or for all users, including e-mail, Lotus Notes, Excel, Text File, Lotus 123W, and the Windows clipboard. If the AllowPrint setting is not specified, the default is to always allow printing.<br><i>Note:</i> This setting affects only existing users who have previously signed on to Analyzer. To apply this setting to new users as they sign on, turn on the DefaultUser option using CHGSCSVRA. The new settings will be returned to the Analyzer Client when the user signs on. It will disable the appropriate commands at that time.                                                                                                                                                                                                                                                                                                                                                              |

Table A-7  
CHGSCSVRA settings for Analyzer functionality (Continued)

| Setting                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AOBlockLoginAfter            | <p>Limits the number of logon attempts for a particular user ID. When the number of consecutive logon failures surpasses the value specified, no further attempts are allowed until a specified time has passed. The default is 3.</p> <p>Note: Logon blocking is based on the client IP and user ID. If a particular user fails to log on, the user is not allowed to log on from the same IP until the expiration period has expired (600 seconds; see BlockLoginFailureRetrySecs below).</p> <p>Upon reaching the point where logins are blocked, the full Windows version of Analyzer will exit. However, Analyzer for the Web does not exit the browser. Nonetheless, the user may disconnect and reconnect at any time.</p> |
| AOBlockLoginFailureRetrySecs | <p>Specifies the time a user must wait before attempting another logon when the failure count is exceeded. This value is in seconds and the default is 600 (10 minutes).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| AOInitialThreadCount         | <p>Defines the initial number of threads created when the Server is started. Since the number of threads will automatically increase if necessary, this setting does not normally need to be specified. Blank for default or number 0 to N. The default is 6 initial threads.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| AOMaximumIdleThreadCount     | <p>Defines the maximum number of idle threads that the server allows. A certain number of idle threads process additional client requests quickly. However, excessive idle threads may use unnecessary system resources. By default, a small number of threads are kept idle. Idle threads in excess of this amount are terminated naturally. This setting does not normally need to be modified. This setting applies to idle threads, <i>not</i> idle client connections. Blank for default or number 0 to N. The default is to have InitialThreadCount + 12 idle threads.</p>                                                                                                                                                  |
| AOMaximumThreadCount         | <p>Defines the maximum number of threads allowed by the Server. Client requests will be processed in the order in which they are received if there are more client requests than threads available to do the work. If this setting is 0, all client requests will be processed on the main thread. This setting takes precedence over the InitialThreadCount and MaximumIdleThreadCount settings. Blank for default or number 0 to N. The default is to have an unlimited amount of threads.</p>                                                                                                                                                                                                                                  |



Table A-7  
CHGSCSVRA settings for Analyzer functionality (Continued)

| Setting            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SecondaryTCPIPPort | <p>Configures the port number that will be associated with the Warehouse Manager Server secondary TCP/IP service. A Warehouse Manager Server can have a secondary port only if it has been upgraded from STRATEGY 3.5 or an earlier release.</p> <p>The service name is formed internally by the Warehouse Manager Server support. It is formed by appending the library name of the Warehouse Manager Server support to SCOLAPSVR. For a Warehouse Manager Server named SCSERVER, the service name becomes SCOLAPSVRSCSERVER. To check for available port numbers, use the OS/400 WRKSRVTBLE command to display existing port assignments.</p> <p>Valid KEYVAL settings can be 5000 to 65535. The default is 23001.</p> <p>When specifying TCPIPPort, a conditional prompt asks whether to update the Analyzer OLAP HTML files with the specified port number. If you are using the Analyzer for the Web JAVA client to connect to the Server through the secondary server TCP/IP port, you must configure the HTML files to use the port the Warehouse Manager Server is listening on.</p> |
| AOTempImagePath    | Define the path on the Server for temporary storage of images used with Analyzer for the Web.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| AOValidateUser     | Controls whether users logging on to the Server from Analyzer or Analyzer for the Web must have a valid iSeries logon. If the setting is 0, an iSeries logon check is performed and the logon is passed directly to Essbase/400. This setting is used if your system has only Essbase/400 profiles and no iSeries profiles. A setting of non-zero causes the Server to validate against an iSeries profile. This is the default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| AOPrimaryLanguage  | Defines the primary language the Server should use for iSeries job logging. This value has no effect on the language used by the client PC connecting to the Server. Valid settings are ENP (uppercase English) and ENU (lowercase English). The default is ENU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| AOClientCCSID      | Specifies the default Windows ANSI code page that clients run under. This setting is used only for the initial communication with the client and if a translation table between the Server CCSID and the client-specified CCSID can't be established. The default client CCSID is 1252.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| TCPIPPort          | <p>Configures the port number associated with TCP/IP service. The number uniquely identifies the TCP/IP service provided by each installation of support. The number specified here creates a TCP/IP service table entry and can't already be defined for any other services. If multiple copies of Warehouse Manager Server support are installed on the iSeries, each copy has its own service name, which is formed by appending the library name of the support to SCSERVER. For a server named SCSERVER, the service name becomes SCSERVERSCAOSVR.</p> <p>Valid settings are from 5000 to 65535. The default is 43419.</p> <p>When specifying TCPIPPort, a conditional prompt asks whether to update the Analyzer HTML files with the specified port number. If you are using the Analyzer for the Web Java Client to connect to AOS, you must configure the HTML files to use the port associated with AOS.</p>                                                                                                                                                                        |

## ***SCCHKPWD Command Parameters***

The SCCHKPWD command searches the local system and lists and/or updates the necessary control table settings with the new password. This ensures uninterrupted scheduled reports and Warehouse Builder distributions.

The following are *required* SCCHKPWD parameters:

USERID (no default) - The name of the user profile to list or update passwords

OLDPWD (no default) - The current password for the USERID

Table A-8 lists the *optional* SCCHKPWD parameters.

Table A-8  
SCCHKPWD *optional parameters*

| Parameter                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UPDATE (default = *NO)    | Indicates whether the passwords for this USERID should be updated in the control tables. Supported values for this parameter are *YES and *NO. Regardless of which option is used, a report of the items matching the specified USERID, OLDPWD, and server(s) is generated and stored in a spool file for the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| DETAILS (default = *YES)  | Indicates whether the generated report lists ALL entries matching the specified criteria or if only a summary of the RDBs, servers, and ODBC data sources affected will be included.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RDB (default = *LOCAL)    | The RDB name of a DB2 database for which passwords may need to be checked or updated. This will affect Warehouse Builder distributions using DB2 sources or targets and scheduled Enterprise Reporting items using ShowCase ODBC data sources to the specified RDB. (Non-STRATEGY data sources must use the ODBC parameter to update passwords.) Special values supported are: *LOCAL, which references the local iSeries database, *ALL, which references all DB2 databases used by items in the catalogs, and *NONE, which indicates that no RDB-related passwords should be checked or updated.                                                                                                                                                                                                                                     |
| AOS (default = *LOCAL)    | The name of a server for which passwords may need to be checked or updated. This will affect only Enterprise Reporting scheduled items using a source or target server. Special values supported are *LOCAL, which references all servers installed on the local system, *ALL, which references any server regardless of the system it's installed on, and *NONE, which indicates that no server-related passwords should be checked or updated.                                                                                                                                                                                                                                                                                                                                                                                       |
| ODBC (default = *NONE)    | The name of an ODBC data source for which passwords may need to be checked or updated. This will affect both Enterprise Reporting scheduled items using DBQs that reference non-STRATEGY data sources and data sources used in WB non-DB2 "pulls." Due to the nature of ODBC, these names are arbitrary, and there is no implicit relationship between the data source name and the underlying system/database (that is, different users may use the same data source name to reference different databases). As a result, take care when using SCCHKPWD to update non-ShowCase ODBC passwords to ensure that the correct data sources are updated. Special values supported are: *ALL, which references all ODBC data sources, and *NONE, which indicates that no (non-STRATEGY) ODBC-related passwords should be checked or updated. |
| ESSBASE (default = *NONE) | The name of an Essbase Server for which a password may need to be checked or updated. Special values supported are: *ALL, which indicates all ESSBASE Servers, and *NONE, which indicates that no ESSBASE Server-related passwords should be checked or updated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| NEWPWD (no default)       | The new password to use for the USERID. This setting is ignored if UPDATE = *NO.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| PWDCONFIRM (no default)   | Confirmation of the new password to use for the USERID (to make sure it was typed correctly). This setting is ignored if UPDATE = *NO.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Essbase/400 Commands

Table A-9 and Table A-10 list common Essbase/400 commands for maintaining an Essbase/400 library.

*Note:* Before using the following commands, add the Essbase library to your iSeries library list and make sure that you have only one Essbase library in your library list. Run the command ADDLIBLE essbaselib, where essbaselib is the name of the Essbase library you want to work with.

Table A-9  
Common Essbase/400 commands

| Command    | Description                                | Usage                                                                                                    |
|------------|--------------------------------------------|----------------------------------------------------------------------------------------------------------|
| CHGESSENV  | Change Essbase/400 environment variables.  | To change the environment.                                                                               |
| DSPESENV   | Display Essbase/400 environment variables. | To display the version or the environment.                                                               |
| DSPESSLOG  | Display Essbase/400 Logs.                  | To read the server and application logs from the iSeries.                                                |
| STRESSSVR  | Start the Essbase/400 Server.              | To start the server and specify a job description to assign to the instance.                             |
| ENDESSSVR  | End the Essbase/400 command processor.     | To stop the server.                                                                                      |
| ESSMSH     | Start MaxL/400 or run MaxL statements.     | To start the MaxL/400 shell or run MaxL statements from a script.                                        |
| GRTESSAUT  | Grant Essbase authority.                   | To allow users lacking proper authority to end/start the Essbase/400 Server.                             |
| RUNMAXLSTM | Run MaxL statements with return codes.     | To run MaxL statements and receive return codes (completion status) that you can use in another program. |
| STRESSCMD  | Start the ESSCMD command processor.        | To run any Essbase/400 command or execute scripts.                                                       |

Table A-10  
Essbase/400 Uninstall command

| Command  | Description                                             | Usage                                                                                                                                                                      |
|----------|---------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNINSTSC | Uninstall Essbase/400 or the Warehouse Manager Library. | To remove an Essbase/400 library from the system. For more information about the UNINSTSC save file, see “Uninstalling a Warehouse Manager Library” on p. 60 in Chapter 5. |

Table A-11 lists common configuration file settings for implementing Integrated Security and specifying how many days of log records to preserve. These keywords are added to the Essbase configuration file. See “Integrated Security” on p. 113 in Chapter 6 for the context in which these keywords are used.

Table A-11  
Common Essbase/400 configuration file settings

| Setting              | Description                                                                | Usage                                                                                                                                                                                                                                 |
|----------------------|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OS400UsersGroup      | Enable all OS/400 users or a subset of OS/400 users to access Essbase/400. | When setting up Integrated Security, this setting gives Essbase/400 access to all users or to a specified group of users.                                                                                                             |
| OS400CommonGroup     | Assign OS/400 user profiles to Essbase/400 groups.                         | When setting up Integrated Security, as OS/400 users are added or removed from OS/400 groups, they will be automatically added or removed from the corresponding Essbase/400 group.                                                   |
| OS400AutoDelete      | Enable automatic deletion of Essbase/400 users.                            | When setting up Integrated Security, this setting automatically deletes all users in the auto-delete group each time the Essbase/400 Server is started.                                                                               |
| OS400DAYSLOGRETAIN   | Specify how many days of Essbase log files to preserve.                    | To set the number of days of log records to keep when the Essbase/400 Server or an Essbase application is started.<br>Example: A value of 2 will retain two days of records.                                                          |
| SHASHCALCWOSHAREDMBR | Improve performance when calculating databases without shared members.     | If CALCOPTCALCHASHTBL is true, insert a SHASHCALCWOSHAREDMBR line in config file specifying each database that doesn't have shared members. Must be applied once per database.<br><b>Example:</b><br>SHASHCALCWOSHAREDMBR<br>tbc demo |



# ***Sample Databases***

This appendix describes the STRATEGY sample databases that you can install with Warehouse Manager Server on the iSeries (AS/400). The databases are:

- SCSAMPLE45, a more summarized database in a star schema format that represents an OLAP-type database
- SCDB45, a fairly normalized database that represents an OLTP-type database

## ***Installing the Sample Databases***

When you install Warehouse Manager Server on the iSeries, you have the option to install the sample database. Selecting this option creates a library on the iSeries called SCSAMPLE45. This library contains the files for the SCSAMPLE45 sample database and the iSeries save files for creating other sample databases, such as SCDB45.

### ***SCSAMPLE45***

The files (tables) for the SCSAMPLE45 sample database are installed in the SCSAMPLE45 library. To install the SCSAMPLE45 library on the iSeries, use one of the following methods:

- Select the option to install the sample library as part of the initial installation of Warehouse Manager Server.
- Run the Restore Sample Library (RSTSMPLB) command on the iSeries command line. For more information, see “Installing the Sample Library after Installation with RSTSMPLB” on p. 162.

### ***SCDB45***

Before you can use the SCDB45 sample database, you must do a small amount of setup on the iSeries. The files for the SCDB45 sample database are shipped in an iSeries save file, called SCDB, within the SCSAMPLE45 library.

*Note:* An iSeries save file is an iSeries object with type \*FILE and attribute \*SAVF. A save file contains data that have been saved using one of the iSeries save commands. To use the data from a save file, you must extract the data from the save file using one of the iSeries restore commands.

To set up the SCDB45 sample database:

- ▶ Install the SCSAMPLE45 database (library).
- ▶ Run the following command on the iSeries command line:

```
RSTLIB SAVLIB(SCDB45) DEV(*SAVF) SAVF(SCSAMPLE45/SCDB) MBROPT(*ALL)
ALWOBJDIF(*ALL) RSTLIB(SCDB45)
```

This command extracts the files (tables) from the SCDB save file in the SCSAMPLE45 library and restores them into the SCDB45 library.

## ***Contents of the Sample Databases***

The sample databases are the basis for the STRATEGY product tutorials and the STRATEGY training manuals. The databases contain data for a fictitious company called The Outdoor Connection. The Outdoor Connection is a sporting goods company that sells through three channels: retail stores, specialty stores, and catalogs. Its primary product lines are skiing equipment, biking equipment, camping gear, and clothing. The Outdoor Connection is based in the United States. It also does business in Canada, Belgium, Germany, France, the United Kingdom, Japan, Australia, and New Zealand.


*Note:* We recommend that you do not modify the original sample database in any way. If it is necessary for you to customize the sample database, copy it to a new library and make your modifications there.

This appendix describes the sample database that you can install with Warehouse Manager Server. The sample database tables are installed into a library called SCSAMPLExx, where xx is the current version of STRATEGY (for example, SCSAMPLE45).

Figure B-1 depicts the general structure of the sample database.



Figure B-1  
Sample database structure



## **SCSAMPLE45 Database**

### **ORDERS Table**

The ORDERS table is the basis of all transaction information. It contains the order number, the location placing the order, the date of the order, the promised delivery date, the date shipped, etc. The table also includes a column that records the location to which each order is shipped. This accommodates the situations in which the order is shipped to a location other than where the order originates.

The ORDERS table also contains the product-related specifics of each order, such as the quantity ordered, quantity shipped, quantity backlogged, and quantity canceled for each product. In addition, the table tracks the unit price, unit cost, extended price, and extended cost for each product ordered. Finally, for international sites, the table includes exchange rates. The rates are used to calculate the foreign extended price and foreign extended cost for each product based on the exchange rate at the time of the order.

The exchange rate is an unsigned percentage and can be mapped to a currency code for each country in which The Outdoor Connection conducts business. It is assumed that all monetary amounts entered into the database are entered in U.S. dollars. Therefore, by definition, the exchange rate for U.S. dollars is 1.0. The exchange rate for other currencies is based on the amount that one U.S. dollar will buy. For example, \$1 U.S. recently bought \$1.4575 CA. Therefore, the exchange rate for Canadian dollars is 1.4575. To convert an amount in the database into local values, multiply the amount by the exchange rate:

$$\$10 \text{ US} \times 1.4575 = \$14.58 \text{ CA} \text{ or } \$10 \text{ US} \times 0.6105 = \text{£}6.11$$

This table contains 14,379 rows of data.

Table B-1  
ORDERS table

| Column     | Description            | Heading              | Data Type     | Contents                                                       |
|------------|------------------------|----------------------|---------------|----------------------------------------------------------------|
| ORDNUM     | Order Number           | Order Number         | Decimal(8)    | Order number                                                   |
| CUST_ID    | Customer ID            | Customer ID          | Numeric(8)    | Account ID of client                                           |
| SHIP_TO    | Shipping Destination   | Ship To              | Numeric(8)    | Account ID of shipping destination                             |
| PRODUCTID  | Product ID (SKU)       | Product ID           | Numeric(8)    | Product ID                                                     |
| PRDLVL03   | Product Level 3        | Product Level 3      | Char(8)       | Third-tier product category                                    |
| QTYORD     | Quantity Ordered       | Quantity Ordered     | Decimal(15)   | Quantity ordered                                               |
| QTYSHP     | Quantity Shipped       | Quantity Shipped     | Decimal(15)   | Quantity shipped                                               |
| QTYBACK    | Quantity Backordered   | Quantity Backordered | Decimal(15)   | Quantity back-ordered                                          |
| QTYCAN     | Quantity Cancelled     | Quantity Cancelled   | Decimal(15)   | Quantity canceled                                              |
| DATEREQ    | Date Requested         | Date Requested       | Date          | Date requested for delivery                                    |
| DATEORD    | Date Ordered           | Date Ordered         | Date          | Date ordered                                                   |
| DATEPROM   | Date Promised          | Date Promised        | Date          | Date promised for delivery                                     |
| DATESHIP   | Date Shipped           | Date Shipped         | Date          | Date shipped                                                   |
| DATEINV    | Date Invoiced          | Date Invoiced        | Date          | Date invoiced                                                  |
| UNIT_PRICE | Unit Price             | Unit Price           | Decimal(15,4) | Unit retail price                                              |
| UNIT_COST  | Unit Cost              | Unit Cost            | Decimal(15,4) | Unit wholesale cost (cost of goods sold)                       |
| EXTPRICE   | Extended Price         | Extended Price       | Decimal(15,4) | Extended retail total (the product of QTYORD and UNIT_PRICE)   |
| EXTCOST    | Extended Cost          | Extended Cost        | Decimal(15,4) | Extended wholesale total (the product of QTYORD and UNIT_COST) |
| MEASURE    | Unit of Measure        | Unit of Measure      | Char(4)       | Unit of measure code                                           |
| CURR_CODE  | Currency Code          | Currency Code        | Char(3)       | Currency code                                                  |
| EXCHG_RATE | Exchange Rate          | Exchange Rate        | Decimal(15,7) | Exchange rate                                                  |
| FRGN_XPRIC | Foreign Extended Price | Foreign Ext Price    | Decimal(15,4) | Foreign extended retail total                                  |
| FRGN_XCOST | Foreign Extended Cost  | Foreign Ext Cost     | Decimal(15,4) | Foreign extended wholesale total                               |
| DIVSN_ID   | Division ID            | Division ID          | Char(8)       | Division code                                                  |
| REGION_ID  | Region ID              | Region ID            | Char(8)       | Region code                                                    |
| AREA_ID    | Area ID                | Area ID              | Char(8)       | Area code                                                      |
| DIST_ID    | Distribution Center ID | Distribution Cntr ID | Char(8)       | Distribution center code                                       |
| REPCODE    | Sales Rep Code         | Rep Code             | Char(4)       | Account representative's code                                  |

**CUSTOMERS Table**

The CUSTOMERS table contains information about The Outdoor Connection's clients. It includes names and addresses and divides the clients into various geographical categories. The Outdoor Connection consists of 3 divisions (North America, Europe, and Asia/Pacific); 12 regions (East, West, South, Midwest, Canada, Belgium, Germany, France, United Kingdom, Japan, Australia, and New Zealand); 70 areas (one for each state and province, including the District of Columbia, plus one for each European and Asia/Pacific country); and 3 distribution facilities (Seattle, Kansas City, and New York). In addition, the CUSTOMERS table includes a currency code for each client, a status code that tracks whether a client is active or inactive, and a code that identifies each client's sales representative.

The CUSTOMERS table contains sufficient data to analyze customers according to two separate hierarchical structures:

Sales organization structure

- Division
- Region
- Area

Geographic location

- Country
- State/Province
- City
- Postal code

This table contains 157 rows of data.

Table B-2  
CUSTOMERS table

| Column      | Description                  | Heading              | Data Type   | Contents                           |
|-------------|------------------------------|----------------------|-------------|------------------------------------|
| CUST_ID     | Customer ID                  | Customer ID          | Numeric(8)  | Account ID of client               |
| CUSTNAME    | Customer Name                | Customer Name        | Char(100)   | Name of client                     |
| LNAME       | Surname                      | Last Name            | Char(50)    | Surname of client contact          |
| FNAME       | Given Name                   | First Name           | Char(50)    | Given name of client contact       |
| ADDRESS1    | Address 1                    | Address 1            | Char(100)   | Address line 1                     |
| ADDRESS2    | Address 2                    | Address 2            | Char(100)   | Address line 2                     |
| CITY        | City                         | City                 | Char(50)    | City name                          |
| STATE       | State/Province Abbreviation  | State/Province Abbr  | Char(2)     | State/province postal abbreviation |
| STATEDESC   | State/Province Description   | State/Province       | Varchar(80) | State/province name                |
| POSTALCODE  | Postal Code                  | Postal Code          | Char(16)    | Postal code                        |
| COUNTY      | County                       | County               | Varchar(80) | County name                        |
| COUNTRY     | Country Code                 | Country Code         | Varchar(12) | Country code                       |
| COUNTRYDESC | Country Description          | Country              | Varchar(80) | Country name                       |
| PHONE       | Phone Number                 | Phone Number         | Char(24)    | Phone number                       |
| E_ADDRESS   | E-mail Address               | E-mail Address       | Char(100)   | E-mail address of client contact   |
| LOB         | LOB                          | LOB Code             | Char(8)     | Line of business code              |
| LOBDESC     | Line of Business Description | Line of Business     | Char(80)    | Line of business description       |
| REPCODE     | Sales Rep Code               | Rep Code             | Char(4)     | Account representative's code      |
| ENTRY_DATE  | Entry Date                   | Entry Date           | Date        | Date added                         |
| ACTIVE      | Active Code                  | Active               | Char(1)     | Active status (Y or N)             |
| DIVSN_ID    | Division ID                  | Division ID          | Char(8)     | Division code                      |
| DIVSN       | Division                     | Division             | Char(80)    | Division name                      |
| REGION_ID   | Region ID                    | Region ID            | Char(8)     | Region code                        |
| REGION      | Region                       | Region               | Char(80)    | Region name                        |
| AREA_ID     | Area ID                      | Area ID              | Char(8)     | Area code                          |
| AREA        | Area                         | Area                 | Char(80)    | Area name                          |
| DIST_ID     | Distribution Center ID       | Distribution Cntr ID | Char(8)     | Distribution center code           |
| DIST_LOC    | Distribution Center          | Distribution Center  | Char(80)    | Distribution center name           |
| CURR_CODE   | Currency Code                | Currency Code        | Char(3)     | Currency code                      |

**PRODUCTS Table**

The PRODUCTS table contains all of The Outdoor Connection's product information and groups products hierarchically. Each product has three levels associated with it. The first identifies a product group (for example, clothing, biking, skiing, etc.), the second, a type (for example, outerwear, shirts, pants, etc.), and the third, a style (for example, black, red, yellow, etc.). In addition, the PRODUCTS table includes a unit price and a unit cost for each product. It also contains a currency code so the unit price and cost can be converted for various countries. Finally, it includes a measures column with a unit type (for example, *one* bicycle, *a pair* of skis, etc.) and a date that records when information was last changed.

This table contains 120 rows of data.

Table B-3  
PRODUCTS table

| Column       | Description           | Heading             | Data Type     | Contents                                 |
|--------------|-----------------------|---------------------|---------------|------------------------------------------|
| PRODUCTID    | Product ID (SKU)      | Product ID          | Numeric(8)    | Product ID                               |
| PRODUCT      | Product Description   | Product Description | Char(80)      | Product description                      |
| PRDLVL01     | Product Level 1       | Product Level 1     | Char(8)       | First tier product category              |
| PRDDESC01    | Product Description 1 | Description Lvl 1   | Char(80)      | First tier product description           |
| PRDLVL02     | Product Level 2       | Product Level 2     | Char(8)       | Second tier product category             |
| PRDDESC02    | Product Description 2 | Description Lvl 2   | Char(80)      | Second tier product description          |
| PRDLV03      | Product Level 3       | Product Level 3     | Char(8)       | Third tier product category              |
| PRDDESC03    | Product Description 3 | Description Lvl 3   | Char(80)      | Third tier product description           |
| PRDLVL04     | Product Level 4       | Product Level 4     | Varchar(8)    | Fourth tier product category             |
| PRDDESC04    | Product Description 4 | Description Lvl 4   | Varchar(27)   | Fourth tier product description          |
| PRDLV05      | Product Level 5       | Product Level 5     | Varchar(8)    | Fifth tier product category              |
| PRDDESC05    | Product Description 5 | Description Lvl 5   | Varchar(27)   | Fifth tier product description           |
| UNITPRICE    | Unit Price            | Unit Price          | Decimal(15,4) | Unit retail price                        |
| UNITCOST     | Unit Cost             | Unit Cost           | Decimal(15,4) | Unit wholesale cost (cost of goods sold) |
| CURR_CODE    | Currency Code         | Currency Code       | Char(3)       | Currency code                            |
| MEASURE      | UOM Code              | UOM Code            | Char(4)       | Unit of measure code                     |
| MEASURE_DESC | Unit of Measure       | Unit of Measure     | Char(80)      | Unit of measure description              |
| ENTRY_DATE   | Entry Date            | Entry Date          | Date          | Entry/update date                        |

**SALESORG Table**

The SALESORG table contains the geographic structure used in the database. Essentially, it maps the various divisions, regions, and areas to each other. It also indicates which distribution facility serves each area or foreign region. As explained previously, the database contains 3 divisions, 12 regions, 70 areas, and 3 distribution facilities.

This table contains 70 rows of data.

Table B-4  
SALESORG table

| Column    | Description            | Heading              | Data Type | Contents                 |
|-----------|------------------------|----------------------|-----------|--------------------------|
| DIVSN_ID  | Division ID            | Division ID          | Char(8)   | Division code            |
| DIVSN     | Division               | Division             | Char(80)  | Division name            |
| REGION_ID | Region ID              | Region ID            | Char(8)   | Region code              |
| REGION    | Region                 | Region               | Char(80)  | Region name              |
| AREA_ID   | Area ID                | Area ID              | Char(8)   | Area code                |
| AREA      | Area                   | Area                 | Char(80)  | Area name                |
| DIST_ID   | Distribution Center ID | Distribution Cntr ID | Char(8)   | Distribution center code |
| DIST_LOC  | Distribution Center    | Distribution Center  | Char(80)  | Distribution center name |

**SALESREPS Table**

The SALESREPS table maps a salesperson's code to a name, employee ID, and manager code. It also maps the salespeople to regions so they can be tracked and evaluated by geography. Each domestic region contains three salespeople; Canada and each European region have two, and each Asia/Pacific region has one.

This table contains 25 rows of data.

Table B-5  
SALESREPS table

| Column     | Description    | Heading      | Data Type    | Contents                      |
|------------|----------------|--------------|--------------|-------------------------------|
| EMP_ID     | Employee ID    | Employee ID  | Numeric(8)   | Employee number               |
| REPCODE    | Sales Rep Code | Rep Code     | Char(4)      | Account representative's code |
| MGRCODE    | Manager Code   | Manager Code | Char(4)      | Supervisor's code             |
| LASTNAME   | Surname        | Last Name    | Varchar(100) | Surname of employee           |
| FIRSTNAME  | Given Name     | First Name   | Varchar(100) | Given name of employee        |
| REGION_ID  | Region ID      | Region ID    | Char(8)      | Region code                   |
| REGION     | Region         | Region       | Varchar(80)  | Region name                   |
| ENTRY_DATE | Entry Date     | Entry Date   | Date         | Entry date (date of hire)     |
| ACTIVE     | Active         | Active       | Char(1)      | Active status (Y or N)        |

**DATEYPES Table**

The DATEYPES table is used to demonstrate how STRATEGY handles date values stored as a data type other than DATE. Specifically, the DATEINV column has been replicated into several formats. You can use this table independently or join it to the ORDERS table on ORDNUM.

This table contains 14,379 rows of data.

Table B-6  
DATEYPES table

| Column           | Description                   | Heading       | Data Type  |
|------------------|-------------------------------|---------------|------------|
| ORDNUM           | Order Number                  | Order Number  | Decimal(8) |
| DATEINV          | Date Invoiced                 | Date Invoiced | Date       |
| DATEINV_YY       | Two Digit Year Value          | YY            | Decimal(2) |
| DATEINV_MM       | Two Digit Month Value         | MM            | Decimal(2) |
| DATEINV_DD       | Two Digit Day Value           | DD            | Decimal(2) |
| DATEINV_EXCELSRL | Five Digit Excel Serial Date  | EXCELSRL      | Decimal(5) |
| DATEINV_YYDDD    | Five Digit YYDDD Date         | YYDDD         | Decimal(5) |
| DATEINV_YYMMDD   | Six Digit YYMMDD Date         | YYMMDD        | Decimal(6) |
| DATEINV_YYDDMM   | Six Digit YYDDMM Date         | YYDDMM        | Decimal(6) |
| DATEINV_DDMMYY   | Six Digit DDMMYY Date         | DDMMYY        | Decimal(6) |
| DATEINV_MMDDYY   | Six Digit MMDDYY Date         | MMDDYY        | Decimal(6) |
| DATEINV_CYYDDD   | Six Digit CYYDDD Date         | CYYDDD        | Decimal(6) |
| DATEINV_CYYMMDD  | Seven Digit CYYMMDD Date      | CYYMMDD       | Decimal(7) |
| DATEINV_CYYDDMM  | Seven Digit CYYDDMM Date      | CYYDDMM       | Decimal(7) |
| DATEINV_YYYYDDD  | Seven Digit YYYYDDD Date      | YYYYDDD       | Decimal(7) |
| DATEINV_YYYYMMDD | Eight Digit YYYYMMDD Date     | YYYYMMDD      | Decimal(8) |
| DATEINV_YYYYDDMM | Eight Digit YYYYDDMM Date     | YYYYDDMM      | Decimal(8) |
| DATEINV_MMDDYYYY | Eight Digit MMDDYYYY Date     | MMDDYYYY      | Decimal(8) |
| DATEINV_DDMMYYYY | Eight Digit DDMMYYYY Date     | DDMMYYYY      | Decimal(8) |
| DATEINV_CHAR     | Eight Character YYYYMMDD Date | YYYYMMDD      | Char(8)    |

**TIMEDIM Table**

The TIMEDIM table contains time-related information for specific dates. This table enables you to perform a variety of time-related analyses.

This table contains 2,192 rows of data.

Table B-7  
TIMEDIM table

| Column    | Description | Heading   | Data Type   | Contents           |
|-----------|-------------|-----------|-------------|--------------------|
| TDATE     | TDATE       | TDATE     | Date        | Date               |
| TYEAR     | TYEAR       | TYEAR     | Integer     | Year               |
| TMONTH    | TMONTH      | TMONTH    | Integer     | Month number       |
| TDAY      | TDAY        | TDAY      | Integer     | Day of month       |
| DAYOFWEEK | DAYOFWEEK   | DAYOFWEEK | Integer     | Day of week number |
| DAYNAME   | DAYNAME     | DAYNAME   | Varchar(9)  | Day                |
| DAYOFYEAR | DAYOFYEAR   | DAYOFYEAR | Integer     | Day of year number |
| TWEEK     | TWEEK       | TWEEK     | Integer     | Week number        |
| HOLCODE   | HOLCODE     | HOLCODE   | Integer     | Holiday code       |
| HOLIDAY   | HOLIDAY     | HOLIDAY   | Varchar(30) | Holiday name       |

**SALESPERF Table**

The SALESPERF table tracks a number of performance measurements. It includes sales projections and actual sales information. It also includes sales quota information for each sales representative and allows you to determine actual and projected sales, as well as quotas, on a geographical basis.

This table contains 12,008 rows of data.

Table B-8  
SALESPERF table

| Column    | Description          | Heading         | Data Type     | Contents                      |
|-----------|----------------------|-----------------|---------------|-------------------------------|
| PERIOD    | Period               | Period          | Date          | Time period                   |
| REPCODE   | Sales Rep Code       | Rep Code        | Char(4)       | Account representative's code |
| AREA_ID   | Area ID              | Area ID         | Char(8)       | Area code                     |
| PRODUCTID | Product ID (SKU)     | Product ID      | Numeric(8)    | Product ID                    |
| PRJ_UNITS | Projected Unit Sales | Projected Units | Numeric(15,0) | Projected unit sales          |
| ACT_UNITS | Actual Unit Sales    | Actual Units    | Numeric(15,0) | Actual units sold             |
| PRJ_SALES | Projected Sales      | Projected Sales | Numeric(15,4) | Projected sales amount        |
| ACT_SALES | Actual Sales         | Actual Sales    | Numeric(15,4) | Actual sales amount           |
| QUOTA     | Quota                | Quota           | Numeric(15,4) | Sales quota                   |



**EMPLOYEES Table**

The EMPLOYEES table contains human resources information about The Outdoor Connection's employees. Specifically, it lists an employee's home address and phone number, e-mail address, gender, date of birth, and, where applicable, social security number. It also contains the employee's compensation type and rate, hire date, and, where applicable, termination date and reason.

This table contains 39 rows of data.

Table B-9  
EMPLOYEES table

| Column     | Description                 | Heading             | Data Type     | Contents                                       |
|------------|-----------------------------|---------------------|---------------|------------------------------------------------|
| EMP_ID     | Employee ID                 | Employee ID         | Numeric(8)    | Employee number                                |
| LASTNAME   | Surname                     | Last Name           | Varchar(100)  | Surname of employee                            |
| FIRSTNAME  | Given Name                  | First Name          | Varchar(100)  | Given name of employee                         |
| ADDRESS1   | Address 1                   | Address 1           | Char(100)     | Address line 1                                 |
| ADDRESS2   | Address 2                   | Address 2           | Char(100)     | Address line 2                                 |
| CITY       | City                        | City                | Char(50)      | City name                                      |
| STATE      | State/Province Abbreviation | State/Province Abbr | Char(2)       | State/province postal abbreviation             |
| POSTALCODE | Postal Code                 | Postal Code         | Char(16)      | Postal code                                    |
| COUNTRY    | Country                     | Country             | Char(50)      | Country name                                   |
| PHONE      | Phone Number                | Phone Number        | Char(24)      | Phone number                                   |
| E_ADDRESS  | E-mail Address              | E-mail Address      | Char(100)     | E-mail address of client contact               |
| SEX        | Gender                      | Sex                 | Char(1)       | Gender                                         |
| DOB        | Date of Birth               | Date of Birth       | Date          | Date of birth                                  |
| SSN        | Social Security Number      | Social Security No. | Char(11)      | Social security number                         |
| COMPTYPE   | Compensation Type           | Comp Type           | Char(1)       | Compensation type (commission, salary, hourly) |
| SALARY     | Base Monthly Salary         | Base Salary         | Numeric(10,2) | Base monthly salary                            |
| HOURLY     | Hourly Rate                 | Hourly Rate         | Numeric(4,2)  | Hourly wage                                    |
| COMMISSION | Commission Rate             | Commission          | Numeric(4,4)  | Commission percentage                          |
| CURR_CODE  | Currency Code               | Currency Code       | Char(3)       | Currency code                                  |
| HIREDATE   | Hire Date                   | Hire Date           | Date          | Hire date                                      |
| TERMDATE   | Termination Date            | Termination Date    | Date          | Termination date                               |
| TERMTYPE   | Termination Reason          | Termination Reason  | Char(20)      | Reason for termination                         |

## SCDB45 Database

The SCDB45 database represents an OLTP-type database. Its tables are mostly normalized, which means that the data are stored in only one place. This requires that tables be joined to find detailed information that may be stored as codes in other tables. For example, the ORDERMAST (Order Master) table contains an order number and a customer ID for each order. To find the details for the order, the ORDERMAST table must be joined to the ORDERDET (Order Detail) table. Likewise, to find information about the customer who placed the order, the ORDERMAST table must be joined to the CUSTMAST (Customer Master) table.

### ORDERMAST Table

The ORDERMAST (Order Master) table is the basis of all transaction information. It contains the “header” information for each order placed, including the order number, the location placing the order, the date of the order, the promised delivery date, the date shipped, etc. The table also includes a column that records the location to which each order is shipped. This accommodates the situations in which the order is shipped to a location other than where the order originates.

This table contains 12,411 rows of data.

Table B-10  
ORDERMAST table

| Column   | Description            | Heading                | Data Type  | Contents                           |
|----------|------------------------|------------------------|------------|------------------------------------|
| ORDNUM   | Order Number           | Order Number           | Decimal(8) | Order number                       |
| CUST_ID  | Customer ID            | Customer ID            | Numeric(8) | Account ID of client               |
| SHIP_TO  | Shipping Destination   | Ship To                | Numeric(8) | Account ID of shipping destination |
| DATEREQ  | Date Requested         | Date Requested         | Date       | Date requested for delivery        |
| DATEORD  | Date Ordered           | Date Ordered           | Date       | Date ordered                       |
| DATEPROM | Date Promised          | Date Promised          | Date       | Date promised for delivery         |
| DATESHIP | Date Shipped           | Date Shipped           | Date       | Date shipped                       |
| DATEINV  | Date Invoiced          | Date Invoiced          | Date       | Date invoiced                      |
| DIST_ID  | Distribution Center ID | Distribution Center ID | Char(8)    | Distribution center code           |
| REPCODE  | Sales Rep Code         | Rep Code               | Char(4)    | Account representative's code      |

**ORDERDET Table**

The ORDERDET (Order Detail) table contains the product-related specifics of each order, such as the quantity ordered, quantity shipped, quantity backlogged, and quantity canceled for each product. In addition, the table tracks the unit price, unit cost, extended price, and extended cost for each product ordered. Finally, for international sites, the table includes exchange rates. The rates are used to calculate the foreign extended price and foreign extended cost for each product based on the exchange rate at the time of the order.

This table contains 14,379 rows of data.

Table B-11  
ORDERDET table

| Column     | Description            | Heading              | Data Type     | Contents                                                       |
|------------|------------------------|----------------------|---------------|----------------------------------------------------------------|
| ORDNUM     | Order Number           | Order Number         | Decimal(8)    | Order number                                                   |
| PRODUCTID  | Product ID (SKU)       | Product ID           | Numeric(8)    | Product ID                                                     |
| PRDLVL03   | Product Level 3        | Product Level 3      | Char(8)       | Third-tier product category                                    |
| QTYORD     | Quantity Ordered       | Quantity Ordered     | Decimal(15)   | Quantity ordered                                               |
| QTYSHP     | Quantity Shipped       | Quantity Shipped     | Decimal(15)   | Quantity shipped                                               |
| QTYBACK    | Quantity Backordered   | Quantity Backordered | Decimal(15)   | Quantity back-ordered                                          |
| QTYCAN     | Quantity Cancelled     | Quantity Cancelled   | Decimal(15)   | Quantity canceled                                              |
| UNIT_PRICE | Unit Price             | Unit Price           | Decimal(15,4) | Unit retail price                                              |
| UNIT_COST  | Unit Cost              | Unit Cost            | Decimal(15,4) | Unit wholesale cost (cost of goods sold)                       |
| EXTPRICE   | Extended Price         | Extended Price       | Decimal(15,4) | Extended retail total (the product of QTYORD and UNIT_PRICE)   |
| EXTCOST    | Extended Cost          | Extended Cost        | Decimal(15,4) | Extended wholesale total (the product of QTYORD and UNIT_COST) |
| MEASURE    | Unit of Measure        | Unit of Measure      | Char(4)       | Unit of measure code                                           |
| CURR_CODE  | Currency Code          | Currency Code        | Char(3)       | Currency code                                                  |
| EXCHG_RATE | Exchange Rate          | Exchange Rate        | Decimal(15,7) | Exchange rate                                                  |
| FRGN_XPRIC | Foreign Extended Price | Foreign Ext Price    | Decimal(15,4) | Foreign extended retail total                                  |
| FRGN_XCOST | Foreign Extended Cost  | Foreign Ext Cost     | Decimal(15,4) | Foreign extended wholesale total                               |

**CUSTMAST Table**

The CUSTMAST (Customer Master) table contains information about The Outdoor Connection's clients. It includes names and addresses and divides the clients into various geographical categories. The Outdoor Connection consists of 3 divisions (North America, Europe, and Asia/Pacific); 12 regions (East, West, South, Midwest, Canada, Belgium, Germany, France, United Kingdom, Japan, Australia, and New Zealand); 70 areas (one for each state and province, including the District of Columbia, plus one for each European and Asia/Pacific country); and 3 distribution facilities (Seattle, Kansas City, and New York). In addition, this table includes a currency code for each client, a status code indicating whether a client is active or inactive, and a code that identifies each client's sales representative.

This table contains 157 rows of data.

Table B-12  
CUSTMAST table

| Column     | Description                 | Heading                     | Data Type   | Contents                           |
|------------|-----------------------------|-----------------------------|-------------|------------------------------------|
| CUST_ID    | Customer ID                 | Customer ID                 | Numeric(8)  | Account ID of client               |
| CUSTNAME   | Customer Name               | Customer Name               | Char(100)   | Name of client                     |
| LNAME      | Surname                     | Last Name                   | Char(50)    | Surname of client contact          |
| FNAME      | Given Name                  | First Name                  | Char(50)    | Given name of client contact       |
| ADDRESS1   | Address 1                   | Address 1                   | Char(100)   | Address line 1                     |
| ADDRESS2   | Address 2                   | Address 2                   | Char(100)   | Address line 2                     |
| CITY       | City                        | City                        | Char(50)    | City name                          |
| STATE      | State/Province Abbreviation | State/Province Abbreviation | Char(2)     | State/province postal abbreviation |
| POSTALCODE | Postal Code                 | Postal Code                 | Char(16)    | Postal code                        |
| COUNTY     | County                      | County                      | Varchar(80) | County name                        |
| COUNTRY    | Country Code                | Country Code                | Varchar(12) | Country code                       |
| PHONE      | Phone Number                | Phone Number                | Char(24)    | Phone number                       |
| E_ADDRESS  | E-mail Address              | E-mail Address              | Char(100)   | E-mail address of client contact   |
| LOB        | LOB                         | LOB Code                    | Char(8)     | Line of business code              |
| REPCODE    | Sales Rep Code              | Rep Code                    | Char(4)     | Account representative's code      |
| ENTRY_DATE | Entry Date                  | Entry Date                  | Date        | Date added                         |
| ACTIVE     | Active Code                 | Active                      | Char(1)     | Active status (Y or N)             |
| DIVSN_ID   | Division ID                 | Division ID                 | Char(8)     | Division code                      |
| REGION_ID  | Region ID                   | Region ID                   | Char(8)     | Region code                        |
| AREA_ID    | Area ID                     | Area ID                     | Char(8)     | Area code                          |
| DIST_ID    | Distribution Center ID      | Distribution Center ID      | Char(8)     | Distribution center code           |
| CURR_CODE  | Currency Code               | Currency Code               | Char(3)     | Currency code                      |

**LOB Table**

The LOB (Line of Business) table maps the line of business code to a description. A line of business can be considered a subclassification for the client. The database contains three lines of business: retail, specialty, and catalog.

This table contains 3 rows of data.

Table B-13

*LOB table*

| Column  | Description                  | Heading          | Data Type | Contents                     |
|---------|------------------------------|------------------|-----------|------------------------------|
| LOB     | LOB                          | LOB Code         | Char(8)   | Line of business code        |
| LOBDESC | Line of Business Description | Line of Business | Char(80)  | Line of business description |

**CURRENCY Table**

The CURRENCY table contains exchange rates used to convert U.S. dollars into foreign currency amounts. The exchange rate is an unsigned percentage and can be mapped to a currency code for each country in which The Outdoor Connection conducts business. It is assumed that all monetary amounts entered into the database are entered in U.S. dollars. Therefore, by definition, the exchange rate for U.S. dollars is 1.0. The exchange rate for other currencies is based on the amount that one U.S. dollar will buy. For example, \$1 U.S. recently bought \$1.4575 CA. Therefore, the exchange rate for Canadian dollars is 1.4575. To convert an amount in the database into local values, multiply the amount by the exchange rate:

$\$10 \text{ US} \times 1.4575 = \$14.58 \text{ CA}$  or  $\$10 \text{ US} \times 0.6105 = \pounds 6.11$

This table contains 9 rows of data.

Table B-14

*CURRENCY table*

| Column     | Description          | Heading       | Data Type     | Contents                        |
|------------|----------------------|---------------|---------------|---------------------------------|
| CURR_CODE  | Currency Code        | Currency Code | Char(3)       | Currency code                   |
| COUNTRY    | Country Code         | Country Code  | Varchar(12)   | Country code                    |
| CURR_DESC  | Currency Description | Currency      | Char(50)      | Currency description            |
| EXCHG_RATE | Exchange Rate        | Exchange Rate | Decimal(15,7) | Exchange rate                   |
| EXCHG_DATE | Effective Date       | Exchange Date | Date          | Effective date of exchange rate |

**CNTRYMAST Table**

The CNTRYMAST (Country Master) table maps the country code to the country name.

This table contains 9 rows of data.

Table B-15  
CNTRYMAST table

| Column      | Description         | Heading      | Data Type   | Contents     |
|-------------|---------------------|--------------|-------------|--------------|
| COUNTRY     | Country Code        | Country Code | Varchar(12) | Country code |
| COUNTRYDESC | Country Description | Country      | Varchar(80) | Country name |

**STATEMAST Table**

The STATEMAST (State Master) table maps the postal abbreviation for the U.S. state or Canadian province to the name of the state or province. The database contains 50 states, 12 provinces, and the District of Columbia.

This table contains 63 rows of data.

Table B-16  
STATEMAST table

| Column    | Description                 | Heading                     | Data Type   | Contents                           |
|-----------|-----------------------------|-----------------------------|-------------|------------------------------------|
| STATE     | State/Province Abbreviation | State/Province Abbreviation | Char(2)     | State/province postal abbreviation |
| STATEDESC | State/Province Description  | State/Province              | Varchar(80) | State/province name                |

**PRODMAST Table**

The PRODMAST (Product Master) table contains all of The Outdoor Connection's product information and groups products hierarchically. Each product has three levels associated with it. The first identifies a product group (for example, clothing, biking, skiing, etc.), the second, a type (for example, outerwear, shirts, pants, etc.), and the third, a style (for example, black, red, yellow, etc.).

In addition, the PRODMAST table includes a unit price and a unit cost for each product. It also contains a currency code so the unit price and cost can be converted for various countries. Finally, it includes a measures column with a unit type (for example, *one* bicycle, *a pair* of skis, etc.) and a date that records when information was last changed.

This table contains 120 rows of data.

Table B-17  
*PRODMAST table*

| Column     | Description         | Heading             | Data Type     | Contents                                 |
|------------|---------------------|---------------------|---------------|------------------------------------------|
| PRODUCTID  | Product ID (SKU)    | Product ID          | Numeric(8)    | Product ID                               |
| PRODUCT    | Product Description | Product Description | Char(80)      | Product description                      |
| PRDLVL01   | Product Level 1     | Product Level 1     | Char(8)       | First tier product category              |
| PRDLVL02   | Product Level 2     | Product Level 2     | Char(8)       | Second tier product category             |
| PRDLV03    | Product Level 3     | Product Level 3     | Char(8)       | Third tier product category              |
| PRDLVL04   | Product Level 4     | Product Level 4     | Varchar(8)    | Fourth tier product category             |
| PRDLV05    | Product Level 5     | Product Level 5     | Varchar(8)    | Fifth tier product category              |
| UNITPRICE  | Unit Price          | Unit Price          | Decimal(15,4) | Unit retail price                        |
| UNITCOST   | Unit Cost           | Unit Cost           | Decimal(15,4) | Unit wholesale cost (cost of goods sold) |
| CURR_CODE  | Currency Code       | Currency Code       | Char(3)       | Currency code                            |
| MEASURE    | UOM Code            | UOM Code            | Char(4)       | Unit of measure code                     |
| ENTRY_DATE | Entry Date          | Entry Date          | Date          | Entry/update date                        |

### ***PLVL01MAST Table***

The PLVL01MAST (Product Level 1 Master) table maps the first product level code to the product level description. This level identifies four product groups: clothing, bicycles, skiing, and camping.

This table contains 4 rows of data.

Table B-18  
*PLVLO1MAST table*

| Column    | Description           | Heading             | Data Type | Contents                       |
|-----------|-----------------------|---------------------|-----------|--------------------------------|
| PRDLVL01  | Product Level 1       | Product Level 1     | Char(8)   | First tier product category    |
| PRDDESC01 | Product Description 1 | Description Level 1 | Char(80)  | First tier product description |

**PLVL02MAST Table**

The PLVL02MAST (Product Level 2 Master) table maps the second product level code to the product level description. This level identifies four product types for each of the product groups. For example, the clothing product group contains the product types outerwear, pants, shirts, and sweaters.

This table contains 16 rows of data.

Table B-19  
PLVL02MAST table

| Column   | Description           | Heading             | Data Type | Contents                        |
|----------|-----------------------|---------------------|-----------|---------------------------------|
| PRDLVL02 | Product Level 2       | Product Level 2     | Char(8)   | Second tier product category    |
| PRDESC02 | Product Description 2 | Description Level 2 | Char(80)  | Second tier product description |

**PLVL03MAST Table**

The PLVL03MAST (Product Level 3 Master) table maps the third product level code to the product level description. This level identifies five product styles: black, blue, green, red, and yellow. The product style can be considered a subclassification for a product.

This table contains 5 rows of data.

Table B-20  
PLVL03MAST table

| Column   | Description           | Heading             | Data Type | Contents                       |
|----------|-----------------------|---------------------|-----------|--------------------------------|
| PRDLV03  | Product Level 3       | Product Level 3     | Char(8)   | Third tier product category    |
| PRDESC03 | Product Description 3 | Description Level 3 | Char(80)  | Third tier product description |

**UOMMAST Table**

The UOMMAST (Unit of Measure Master) table maps the unit of measure code to the measure description. The unit of measure identifies whether a product is sold as a single unit or in pairs.

This table contains 2 rows of data.

Table B-21  
UOMMAST table

| Column       | Description     | Heading         | Data Type | Contents                    |
|--------------|-----------------|-----------------|-----------|-----------------------------|
| MEASURE      | UOM Code        | UOM Code        | Char(4)   | Unit of measure code        |
| MEASURE_DESC | Unit of Measure | Unit of Measure | Char(80)  | Unit of measure description |



**ORGMAS<sup>T</sup> Table**

The ORGMAS<sup>T</sup> (Organization Master) table contains the geographic structure used in the database. Essentially, it maps the various divisions, regions, and areas to each other. Each division has several regions, and each region has several areas. This table also indicates which distribution facility serves each area or foreign region. The database contains 3 divisions, 12 regions, 70 areas, and 3 distribution facilities.

This table contains 70 rows of data.

Table B-22  
ORGMAS<sup>T</sup> table

| Column    | Description            | Heading                | Data Type | Contents                 |
|-----------|------------------------|------------------------|-----------|--------------------------|
| DIVSN_ID  | Division ID            | Division ID            | Char(8)   | Division code            |
| REGION_ID | Region ID              | Region ID              | Char(8)   | Region code              |
| AREA_ID   | Area ID                | Area ID                | Char(8)   | Area code                |
| DIST_ID   | Distribution Center ID | Distribution Center ID | Char(8)   | Distribution center code |

**DIVSNMAS<sup>T</sup> Table**

The DIVSNMAS<sup>T</sup> (Division Master) table maps each division code to a division name. The database contains three divisions: North America, Europe, and Asia/Pacific.

This table contains 3 rows of data.

Table B-23  
DIVSNMAS<sup>T</sup> table

| Column   | Description | Heading     | Data Type | Contents      |
|----------|-------------|-------------|-----------|---------------|
| DIVSN_ID | Division ID | Division ID | Char(8)   | Division code |
| DIVSN    | Division    | Division    | Char(80)  | Division name |

**REGIONMAS<sup>T</sup> Table**

The REGIONMAS<sup>T</sup> (Region Master) table maps each region code to a region name. The database contains three North American regions, four European regions, and three Asia/Pacific regions.

This table contains 12 rows of data.

Table B-24  
REGIONMAS<sup>T</sup> table

| Column    | Description | Heading   | Data Type | Contents    |
|-----------|-------------|-----------|-----------|-------------|
| REGION_ID | Region ID   | Region ID | Char(8)   | Region code |
| REGION    | Region      | Region    | Char(80)  | Region name |

**AREAMAST Table**

The AREAMAST (Area Master) table maps each area code to an area name. The database contains 63 areas in the North American region, 4 areas in the European region, and 3 areas in the Asia/Pacific region.

This table contains 70 rows of data.

Table B-25

*AREAMAST table*

| Column  | Description | Heading | Data Type | Contents  |
|---------|-------------|---------|-----------|-----------|
| AREA_ID | Area ID     | Area ID | Char(8)   | Area code |
| AREA    | Area        | Area    | Char(80)  | Area name |

**DISTMAST Table**

The DISTMAST (Distribution Center Master) table maps each distribution center code to a distribution center name. The database contains three distribution centers: Seattle, Kansas City, and New York. Each distribution center can serve all divisions.

This table contains 3 rows of data.

Table B-26

*DISTMAST table*

| Column   | Description            | Heading                | Data Type | Contents                 |
|----------|------------------------|------------------------|-----------|--------------------------|
| DIST_ID  | Distribution Center ID | Distribution Center ID | Char(8)   | Distribution center code |
| DIST_LOC | Distribution Center    | Distribution Center    | Char(80)  | Distribution center name |

**DATEYPES Table**

The DATEYPES table is used to demonstrate how STRATEGY handles date values stored as a data type other than DATE. Specifically, the DATEINV column has been replicated into several formats. You can use this table independently or join it to the ORDERMAST table on ORDNUM.

This table contains 12,411 rows of data.

Table B-27  
DATEYPES table

| Column           | Description                                | Heading       | Data Type  |
|------------------|--------------------------------------------|---------------|------------|
| ORDNUM           | Order Number                               | Order Number  | Decimal(8) |
| DATEINV          | Date Invoiced                              | Date Invoiced | Date       |
| DATEINV_YY       | Two Digit Year Value                       | YY            | Decimal(2) |
| DATEINV_MM       | Two Digit Month Value                      | MM            | Decimal(2) |
| DATEINV_DD       | Two Digit Day Value                        | DD            | Decimal(2) |
| DATEINV_EXCELSRL | Five Digit Excel Serial Date               | EXCELSRL      | Decimal(5) |
| DATEINV_YYDDD    | Five Digit YYDDD Date                      | YYDDD         | Decimal(5) |
| DATEINV_YYMMDD   | Six Digit YYMMDD Date                      | YYMMDD        | Decimal(6) |
| DATEINV_YYDDMM   | Six Digit YYDDMM Date                      | YYDDMM        | Decimal(6) |
| DATEINV_DDMMYY   | Six Digit DDMMYY Date                      | DDMMYY        | Decimal(6) |
| DATEINV_MMDDYY   | Six Digit MMDDYY Date                      | MMDDYY        | Decimal(6) |
| DATEINV_CYYDDD   | Six Digit CYYDDD Date                      | CYYDDD        | Decimal(6) |
| DATEINV_CYYMMDD  | Seven Digit CYYMMDD Date                   | CYYMMDD       | Decimal(7) |
| DATEINV_CYYDDMM  | Seven Digit CYYDDMM Date                   | CYYDDMM       | Decimal(7) |
| DATEINV_YYYYDDD  | Seven Digit YYYYDDD Date                   | YYYYDDD       | Decimal(7) |
| DATEINV_YYYYMMDD | Eight Digit YYYYMMDD Date                  | YYYYMMDD      | Decimal(8) |
| DATEINV_YYYYDDMM | Eight Digit YYYYDDMM Date                  | YYYYDDMM      | Decimal(8) |
| DATEINV_MMDDYYYY | Eight Digit MMDDYYYY Date                  | MMDDYYYY      | Decimal(8) |
| DATEINV_DDMMYYYY | Eight Digit DDMMYYYY Date                  | DDMMYYYY      | Decimal(8) |
| DATEINV_CHAR     | Eight Character YYYYMMDD Date              | YYYYMMDD      | Char(8)    |
| DATEINV_HYF      | Five Digit Infinium Date<br>(EXCELSRL - 1) | HYF           | Decimal(5) |

**SALESPERF Table**

The SALESPERF table tracks a number of performance measurements. It includes sales projections and actual sales information. It also includes sales quota information for each sales representative and allows you to determine actual and projected sales, as well as quotas, on a geographical basis.

This table contains 12,008 rows of data.

Table B-28  
SALESPERF table

| Column    | Description          | Heading         | Data Type     | Contents                      |
|-----------|----------------------|-----------------|---------------|-------------------------------|
| PERIOD    | Period               | Period          | Date          | Time period                   |
| REPCODE   | Sales Rep Code       | Rep Code        | Char(4)       | Account representative's code |
| AREA_ID   | Area ID              | Area ID         | Char(8)       | Area code                     |
| PRODUCTID | Product ID (SKU)     | Product ID      | Numeric(8)    | Product ID                    |
| PRJ_UNITS | Projected Unit Sales | Projected Units | Numeric(15,0) | Projected unit sales          |
| ACT_UNITS | Actual Unit Sales    | Actual Units    | Numeric(15,0) | Actual units sold             |
| PRJ_SALES | Projected Sales      | Projected Sales | Numeric(15,4) | Projected sales amount        |
| ACT_SALES | Actual Sales         | Actual Sales    | Numeric(15,4) | Actual sales amount           |
| QUOTA     | Quota                | Quota           | Numeric(15,4) | Sales quota                   |

**EMPMASST Table**

The EMPMASST (Employee Master) table contains human resources information about The Outdoor Connection's employees. Specifically, it lists an employee's home address and phone number, e-mail address, gender, date of birth, and, where applicable, social security number. It also contains the employee's compensation type and rate, hire date, and, where applicable, termination date and reason. If an employee is part of the sales organization, the table contains the employee's sales representative code, the representative code of the employee's manager, the date the employee entered the sales organization, and a status code indicating whether the employee is currently active or inactive in the sales organization.

This table contains 39 rows of data.

Table B-29  
EMPMAST table

| Column     | Description                 | Heading                     | Data Type     | Contents                                       |
|------------|-----------------------------|-----------------------------|---------------|------------------------------------------------|
| EMP_ID     | Employee ID                 | Employee ID                 | Numeric(8)    | Employee number                                |
| LASTNAME   | Surname                     | Last Name                   | Varchar(100)  | Surname of employee                            |
| FIRSTNAME  | Given Name                  | First Name                  | Varchar(100)  | Given name of employee                         |
| ADDRESS1   | Address 1                   | Address 1                   | Char(100)     | Address line 1                                 |
| ADDRESS2   | Address 2                   | Address 2                   | Char(100)     | Address line 2                                 |
| CITY       | City                        | City                        | Char(50)      | City name                                      |
| STATE      | State/Province Abbreviation | State/Province Abbreviation | Char(2)       | State/province postal abbreviation             |
| POSTALCODE | Postal Code                 | Postal Code                 | Char(16)      | Postal code                                    |
| COUNTRY    | Country                     | Country                     | Varchar(12)   | Country code                                   |
| PHONE      | Phone Number                | Phone Number                | Char(24)      | Phone number                                   |
| E_ADDRESS  | E-mail Address              | E-mail Address              | Char(100)     | E-mail address of employee                     |
| SEX        | Gender                      | Sex                         | Char(1)       | Gender                                         |
| DOB        | Date of Birth               | Date of Birth               | Date          | Date of birth                                  |
| SSN        | Social Security Number      | Social Security No.         | Char(11)      | Social security number                         |
| COMPTYPE   | Compensation Type           | Comp Type                   | Char(1)       | Compensation type (commission, salary, hourly) |
| SALARY     | Base Monthly Salary         | Base Salary                 | Numeric(10,2) | Base monthly salary                            |
| HOURLY     | Hourly Rate                 | Hourly Rate                 | Numeric(4,2)  | Hourly wage                                    |
| COMMISSION | Commission Rate             | Commission                  | Numeric(4,4)  | Commission percentage                          |
| CURR_CODE  | Currency Code               | Currency Code               | Char(3)       | Currency code                                  |
| HIREDATE   | Hire Date                   | Hire Date                   | Date          | Hire date                                      |
| TERMDATE   | Termination Date            | Termination Date            | Date          | Termination date                               |
| TERMTYPE   | Termination Reason          | Termination Reason          | Char(20)      | Reason for termination                         |
| REPCODE    | Sales Rep Code              | Rep Code                    | Char(4)       | Sales rep code                                 |
| MGRCODE    | Manager Code                | Manager Code                | Char(4)       | Sales manager code                             |
| ENTRY_DATE | Entry Date as Sales Rep     | Entry Date                  | Date          | Entry date into sales organization             |
| ACTIVE     | Active Sales Rep            | Active                      | Char(1)       | Active status of sales rep                     |

## ***Installing the Sample Library after Installation with RSTSMPLB***

RSTSMPLB (Restore Sample Library) will install the sample databases if they were not previously installed. This command can also restore the sample database back to its original structure (if database files were altered or destroyed). To use this command, type RSTSMPLB at the command prompt and press Enter.

The command has two parameters:

- CLRLIB (\*YES, \*NO). If \*YES is chosen, the sample database library (if it exists) is cleared (all files within destroyed) before the STRATEGY sample files are restored. Sample files are restored to their original version installed with this release.
- PRVLGL (\*YES, \*NO). If \*YES is chosen, all logical files on the system that are dependent on the STRATEGY sample files within the sample database library are preserved. Logical files that are dependent on non-STRATEGY files within the sample library are not preserved.

*Note:* If the user wants to preserve logical files not dependent on STRATEGY sample files, choose \*NO for CLRLIB.

- access
  - path, 28
  - plan, 28
- adding
  - databases, 84
- advantages of
  - run limits, 18
  - run priorities, 19
- ALLOBJ authority, 113
- allow parameterization, 31
- Analyzer
  - default user profiles, 85
- Analyzer for the Web
  - adding databases, 84
  - APPLET tag components, 86
  - firewall and proxy server support, 79
  - iSeries (AS/400) profiles, 85
  - symbolic link, 81
  - tailoring Web pages, 85
- Analyzer OLAP functions
  - changing the ASCII CCSID, 83
  - changing the EBCDIC CCSID, 82
  - configuring the iSeries (AS/400) HTTP server, 75
  - displaying the ASCII CCSID, 83
  - displaying the current EBCDIC CCSID, 82
  - displaying the EBCDIC CCSID, 82
  - finding connections, 59
  - HTML documents, 83
- analyzing queries, 29
- APPLET tag components, 86
- AREAMAST table, 158
- AS/400
  - command reference, 127
  - installing STRATEGY, iv
  - managing STRATEGY data, 41
  - profiles, 85
  - work management, 68
- AS/400 commands, 127
  - default security settings, 71
  - reference, 127
  - security issues, 71
- ASCII CCSID, 97
- Asia Pacific support, 3
- assigning
  - job run priority, 20
  - levels of parallelism, 21
  - query run limits, 19
- auditing
  - additional analysis tips, 26
  - analyzing the information, 24
  - collecting the information, 23
  - impact to overall performance, 23
  - managing your environment, 27
  - to understand queries, 23
- authority
  - required for installation, 5, 6
- automatic deletion of Essbase/400 users, 122
- auxiliary storage, 36
- avoiding
  - pattern matching, 32
  - string padding, 32
- backing up
  - a Warehouse Manager library, 61
  - data, 41
  - Essbase/400 files, 124
  - server information, 61
  - STRATEGY data, 43
- batch queries, 69, 70
- CAB files
  - uninstalling, 88
- changed records, 34
  - sending during distribution, 34
- changing
  - data views with dependent data views, 47
  - ODBC.INI settings, 31
  - physical files with dependent data views, 47
- checking current usage
  - relational products, 58, 59
- CHGJOB CL, 20
- CHGPF command, 63
- CHGSCSVRA, 128
- CHGSRVRSTS command, 66
- CLASSPATH, 94
- CLRPFM, 34
- CMPSCJAR command, 71, 130
- CNTRYMAST table, 154
- command line
  - MaxL/400, 101

- commands
  - default security settings, 71
  - SBMDD, 47, 50
  - SCCHKPWD, 48, 130
  - security issues, 71
  - STRDD, 47, 50
- compile java files, 130
- configuration after restoring server, 46
- configuration file
  - editing *ShowCase.cfg* for J.D. Edwards, 53
  - Essbase/400, 92
- configuring
  - a new HTTP server for Analyzer for the Web, 75
  - an existing HTTP server for Analyzer for the Web, 81
  - Essbase/400 environment parameters, 94
  - Essbase/400 information, 91
  - the Essbase environment, 94
  - iSeries (AS/400) HTTP Server, 75
  - Orion Server for Analyzer for the Web, 78
  - Warehouse Manager TCP/IP Support, 57
- connections
  - Analyzer OLAP, 59
- contacting
  - documentation team, iv
  - support, 1
- control tables, 48
- controlling
  - level of parallelism, 21
  - Query/Report Writer system usage, 18
- CURRENCY table, 153
- CUSTOMERS table, 143
- CUSTOMERS table, 152
- custom library lists for users, 43
- CUSTOMERS table, 143
- CUSTOMERS table, 152
- data
  - backing up, 43
  - disaster recovery backups, 44
  - loading into Essbase/400, 109, 111
  - reloading rejected records, 112
  - temporary backups, 44
- data load
  - rejected records, 112
- data views
  - changing dependent data views, 47
  - changing physical files, 47
- DATETYPES table, 147, 159
- DB2 symmetric multiprocessing, 23
- DB2/400
  - loading data into Essbase/400, 109
- debugging, 108
- default wait time, 68
- defining the uses of an Essbase/400 application, 38
  - deleting target tables upon distribution, 34
  - detecting and removing obsolete information, 65
  - determining level of parallelism, 21
  - dimension build
    - rejected records, 112
  - directing data, 42
  - disadvantages of
    - run limits, 18
    - run priorities, 19
  - disaster recovery backups, 44
  - displaying
    - Essbase/400 environment parameters, 95
  - DISTMAST table, 158
  - distributed archive files, 88
  - distributed relational database architecture, 74
  - distributions
    - choosing the most efficient, 33
  - DIVSNMAST table, 157
  - documentation
    - online, iv
    - providing direct feedback, iv
  - DRDA, 74
- EBCDIC CCSID, 96
- editing *QSYS.LIB* files from a 5250 display session, 93
- EDTF, 127
- EDTF command, 93
- e-mail support, 1
- EMPLOYEES table, 149
- EMPMASST table, 160
- enabling expert cache, 22
- ending
  - TCP/IP, 60
  - Warehouse Manager TCP/IP, 60
- ENDSCSVR, 128
- enhancement requests, 171
- Enterprise Reporting
  - backing up data, 46
- Essbase NT, 50
- ESSBASE.CFG* file, 92
- Essbase/400
  - ASCII CCSID, 97
  - backing up files, 124
  - calculation process, 37
  - command reference, 136
  - commands, 136
  - configuration and environment relationship, 95
  - configuration file, 92
  - configuring server information, 91
  - configuring the Essbase environment, 94
  - debugging, 108



- defining uses, 38
- EBCDIC CCSID, 96
- ESSCMD/400, 98
- index cache size, 109
- installation library and path, 91
- Integrated Security, 114
- iSeries (AS/400) performance tuning, 37
- Java, 95
- jobs, 123
- loading data, 109
- loading databases from Warehouse Builder, 111
- log files, 108
- maintenance, 91
- message logging, 124
- national language support, 95
- passwords, 123
- performance, 36
- preventing storage overflow, 36
- QSYS.LIB* objects, 91, 93
- QSYS.LIB* objects created during installation, 91
- reloading rejected records, 112
- removing Essbase environment parameters, 95
- restoring files, 125
- script files, 98, 101
- security, 112
- special considerations, 91
- SQL data load plan, 36
- Essbase/400 Server
  - changing the ASCII and EBCDIC CCSIDs, 97
  - displaying the current ASCII CCSID, 97
  - displaying the current EBCDIC CCSID, 96
  - starting, 5, 6
  - stopping, 7
- ESSCMD API Toolkit, 103
- ESSCMD/400, 98
  - script file differences, 101
- ESSMSH command, 99
- European support, 3
- evaluating queries with performance analyzer, 29
- expert cache
  - enabling, 22
- fax support, 1
- files
  - configuration and log, 92
- finding
  - available port numbers, 58
- firewall and proxy servers
  - Analyzer for the Web, 79
  - configuring AOSServlet with IBM Websphere Server, 79
- flat file support, 111
- getting help on ESSCMD commands, 103
- HTTP server
  - configuring, 75
- improving
  - performance with indexes, 30
  - Query performance, 28
- index
  - cache size, 109
  - creating, 30
- installing
  - iSeries (AS/400) media and passwords, 63
  - Orion Server, 7
  - STRATEGY, iv
- installing the sample databases with RSTSMPLB, 162
- Integrated Security, 113, 114
- Internet support, 1
- iSeries (AS/400)
  - installing STRATEGY, iv
  - managing STRATEGY data, 41
  - profiles, 85
  - work management, 68
- iSeries (AS/400) commands, 127
  - default security settings, 71
  - reference, 127
- J.D. Edwards
  - editing *ShowCase.cfg* to use alternate F0094 file, 53
  - group library list selection order for OneWorld, 54
  - group library list selection order for World, 54
  - integration with STRATEGY, 52
  - multiple library list environments, 54
  - non-iSeries profile sign-on, 52
  - user-defined codes (UDC), 54
- Japan and Asia Pacific support, 3
- JAR files, 89
- Java
  - CMPCJAR command, 130
- Java and Essbase/400, 95
- job
  - class, 68
  - description, 70
  - run priority, 68
- job run priorities
  - setting, 19
- journals
  - working with, 67
- JVMModuleLocation parameter, 95

- library
  - backing up a server, 61
  - checking current usage, 58
  - migrating to a new, 58
  - releasing the restriction on a server, 67
- library lists
  - creating for users, 43
- licensing, iv
- limiting
  - data type conversion, 32
  - use of nested queries, 32
- loading
  - data into Essbase/400, 109, 111
  - DB2/400 data into Essbase/400, 109
- LOB table, 153
- log files, 108
- LOGINESS command, 104
- LOGOUTESS marker, 104
  
- maintaining Essbase/400 log files, 108
- managing
  - STRATEGY data on the iSeries (AS/400), 41
  - your auditing environment, 27
- maximum temporary storage, 68
- MaxL/400, 98
  - command line interface, 101
  - ESSMSH command, 99
  - input script file, 98
  - RUNMAXLSTM command, 101
  - scripts, 99
- message logging, 124
- migrating
  - data from test to production, 41
  - server information, 62
  - server information for backup purposes, 43
  - to a new library, 58
  - users to a new Warehouse Manager library, 58
- minimizing the use of LIKE, 32
- miscellaneous objects, 92
- moving
  - data from test to production, 41
  - server information, 62
  - server information for backup purposes, 43
- multiprocessing, 23
  
- national language support
  - Essbase/400, 95
- Netscape Communicator and JAR files, 89
- North American support, 2
  
- obsolete information
  - removing, 65
- ODBC.INI* settings
  - changing, 31
- online documentation, iv
- optimization level, 31
- optimizing
  - Warehouse Builder performance, 32
- ORDERDET table, 151
- ORDERMAST table, 150
- ORDERS table, 141
- ORGMAST table, 157
- Orion Server
  - installing, 7
  - starting, 7
  - stopping, 7
- OS400DAYSLOGRETAIN, 108, 137
- OUTPUT command, 102
  
- parallelism, 21
  - assigning levels, 21
- parameterization, 31
- passwords, iv
  - Essbase/400, 123
  - managing, 48
- pattern matching
  - avoiding, 32
- performance
  - system tuning, 22
- performance analyzer, 19
  - evaluating queries, 29
  - using, 29
- PLVL01MAST table, 155
- PLVL02MAST table, 156
- PLVL03MAST table, 156
- pre-creating libraries, 34
- PRODMAST table, 154
- product enhancement request form, 171
- production environment
  - moving from a test environment, 41
- PRODUCTS table, 145
- pull distributions, 33
  
- QSYS.LIB* objects, 91, 93
- query creation tips, 32
- query optimizer
  - overview, 28
- Query performance
  - improving, 28

- Query run limits
  - setting maximum limits, 18
- redirecting STRATEGY data, 42
- REGIONMAST table, 157
- rejected records
  - reloading, 112
- releasing a Warehouse Manager library
  - from a restricted state, 67
- reloading rejected records, 112
- removing
  - a Warehouse Manager library, 60
  - Essbase/400 environment parameters, 95
  - obsolete information, 65
  - Warehouse Manager Server libraries, 60
  - Warehouse Manager Server libraries (pre-STRATEGY), 61
- renaming users and groups, 123
- resource manager, 18
- Restore Sample Library, 162
- restoring
  - Essbase/400 files, 125
  - server information, 63
- restoring server
  - follow-up configuration, 46
- restricting access to a Warehouse Manager library, 65
- routing entries, 69
- RSTSMPLB command, 162
- RTVESSSTS command, 104
- run limits
  - advantages, 18
  - assigning, 19
  - disadvantages, 18
  - indexing, 18
  - verifying, 19
- run priorities
  - advantages, 19
  - assigning, 20
  - disadvantages, 19
  - setting, 19
- RUNESSCMD command, 104
- RUNMAXLSTM command, 101
- SALESORG table, 146
- SALESPERF table, 148, 160
- SALESREPS table, 146
- sample databases, 139
  - installing, 162
- save and restore, 41
- Save and Restore feature, 62
- saving
  - (AS/400) media and passwords, 61
  - server information, 61
- SBMDD command, 47, 50, 130
- SCCHKPWD command, 48, 130
  - parameters, 134
- SCDB45, 139
- SCDB45 sample database
  - installing, 140
- scenarios
  - test and production on separate iSeries (AS/400s), 42
  - test and production on the same iSeries (AS/400), 43
- SCHASHCALCWOSHAREDMBR setting, 137
- script files, 98, 101
- SCSAMPLE45, 139
- SCSAMPLE45 sample database
  - installing, 139
- security, 112
  - Essbase/400, 112
  - iSeries (AS/400) command authorities, 71
  - issues, 71
- sequence numbers
  - using, 35
- server information
  - backing up (saving), 61
  - restoring, 63
- servers
  - backing up, 43
  - releasing the restriction on, 67
  - starting, 5
  - stopping, 6, 7
- Setting, 18
- setting
  - Essbase/400 Integrated Security, 114
  - maximum Query run limits, 18
  - run limits, 18
  - server job run priorities, 19
- shared member, 137
- specifying
  - join conditions, 32
  - output and error files within ESSCMD/400 script files, 100
- SQL
  - drill-through using Essbase add-ins, 124
  - functions to control distributions, 35
  - loading data into Essbase/400, 109
  - optimization level, 31
- starting
  - Essbase/400 Server, 5
  - Orion Server, 7
  - Warehouse Manager Server, 5
- starting and stopping the servers, 5
- STATEMAST table, 154

- stopping
  - Essbase/400 Server, 7
  - Orion Server, 7
  - Warehouse Manager Server, 6
- storage pools, 69
- STRATEGY
  - migrating to a new library, 58
  - passwords and licensing, iv
  - security issues, 71
- STRDD command, 47, 50, 130
- STRESSCMD command, 99
- string matching
  - avoiding, 32
- string padding
  - avoiding, 32
- STRSCSVR, 129
- STRSCSVRA, 57
- substitution variables in rule file SQL statements, 110
- subsystem
  - Essbase/400 Server, 6
- subsystems, 69
- support
  - Internet, 1
- symbolic link, 81
- system performance tuning, 22
- system usage
  - controlling, 18
- tailoring Analyzer Web pages, 85
- target tables
  - deleting upon distribution, 34
- TCP/IP, 57, 69
  - ending, 60
  - job descriptions, 70
  - port numbers, 58
- telephone support, 1
- temporary backups, 44
- test environment
  - moving to production, 41
- The Outdoor Connection, 140
- TIMEDIM table, 148
- tips
  - always specify join conditions, 32
  - avoid string padding, 32
  - avoid use of arithmetical expressions in WHERE clause, 32
  - limit use of data type conversion, 32
  - limit use of nested queries, 32
  - minimize use of LIKE, 32
  - query creation, 32
- tuning system performance, 22
- types of distributions, 33
- understanding
  - APPLET tag components, 86
  - configuration and environment relationship, 95
  - Essbase/400 configuration, 91
  - query optimizer, 28
- uninstalling
  - a Warehouse Manager library, 60
  - CAB files, 88
  - Warehouse Manager Server libraries, 60
  - Warehouse Manager Server libraries (pre-STRATEGY), 61
- UOMMAST table, 156
- user
  - default profiles for Analyzer, 85
  - exit programs, 73
  - profiles, 70
- user exit programs, 73
- using
  - Essbase/400 configuration file settings, 92
  - Essbase/400 environment settings, 93
  - Essbase/400 national language support (NLS), 95
  - ESSCMD API Toolkit, 103
  - QSYS.LIB files from a 5250 display session, 93
  - TCP/IP with STRATEGY, 57
  - Warehouse Builder with Essbase/400, 111
- Using CURDATE (), 35
- verifying the run limit, 19
- Warehouse Builder
  - improving performance, 32
  - naming the Essbase/400 Server, 112
  - tuning issues, 35
  - using to load Essbase databases, 111
- Warehouse Manager
  - backing up a server library, 61
  - checking current usage, 58
  - command reference, 128
  - commands, 128
  - ending TCP/IP, 60
  - releasing the restriction on a server library, 67
  - removing obsolete information, 65
  - security issues, 71
  - work management, 68
  - working with journals, 67
- Warehouse Manager and iSeries (AS/400)
  - work management, 68

- Warehouse Manager commands
  - authorities, 71
  - default security settings, 71
- Warehouse Manager Server
  - starting, 5
  - stopping, 6
- WHERE clause
  - avoiding arithmetical expressions, 32
- work management, 68
- work management controls, 68
  - job class, 68
  - job description, 70
  - routing entries, 69
  - storage pool, 69
  - subsystems, 69
  - user profile, 70
- working with journals, 67



# Product Enhancement Request Form

Use this fax-back form to send ShowCase, a division of SPSS Inc., a description of enhancements to STRATEGY that you would like included in future releases. To help us prioritize requests, please include a description of how the change will benefit your business. Attach additional sheets if necessary.

| Product Enhancement Request                                                                                                                                                                                                        |                                                                                                                                                             |                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| <b>Return To:</b><br><br>Kathy Maegerlein<br>ShowCase, a division of SPSS Inc.<br>4115 Highway 52 North, Suite 300<br>Rochester, MN 55901-0144<br>Telephone: (507) 288-5922<br>Fax: (507) 287-2803<br>E-mail: kmaegerlein@spss.com | <b>From:</b><br><br>Date:<br><br>Title:<br><br>Company:<br><br>Country:<br><br>Telephone:<br><br>Fax:<br><br>Pages, including cover:<br><br>E-mail address: |                                               |
| STRATEGY Product(s) (check those that apply)                                                                                                                                                                                       |                                                                                                                                                             |                                               |
| <input type="checkbox"/> Warehouse Builder                                                                                                                                                                                         | <input type="checkbox"/> Essbase/400                                                                                                                        | <input type="checkbox"/> Analyzer for the Web |
| <input type="checkbox"/> Warehouse Manager (server)                                                                                                                                                                                | <input type="checkbox"/> Query                                                                                                                              | <input type="checkbox"/> Documentation        |
| <input type="checkbox"/> Warehouse Manager (client)                                                                                                                                                                                | <input type="checkbox"/> Report Writer                                                                                                                      | <input type="checkbox"/> Other:               |
| <input type="checkbox"/> Tutorials/Samples                                                                                                                                                                                         | <input type="checkbox"/> Analyzer                                                                                                                           |                                               |
| Enhancement Description                                                                                                                                                                                                            |                                                                                                                                                             |                                               |
| <br><br><br><br><br><br><br><br><br><br>                                                                                                                                                                                           |                                                                                                                                                             |                                               |
| Benefit to Your Business                                                                                                                                                                                                           |                                                                                                                                                             |                                               |
| <br><br><br><br><br><br><br><br><br><br>                                                                                                                                                                                           |                                                                                                                                                             |                                               |

