



Administrator Guide Showcase 9



Copyright Terms and Conditions

The content in this document is protected by the Copyright Laws of the United States of America and other countries worldwide. The unauthorized use and/or duplication of this material without express and written permission from HelpSystems is strictly prohibited. Excerpts and links may be used, provided that full and clear credit is given to HelpSystems with appropriate and specific direction to the original content. HelpSystems and its trademarks are properties of the HelpSystems group of companies. All other marks are property of their respective owners.

201803095145

About HelpSystems

HelpSystems is a leading provider of systems & network management, business intelligence, and security & compliance software. We help businesses reduce data center costs by improving operational control and delivery of IT services.

Preface

Welcome to the Showcase 9 Administrator's Guide. This guide provides information about the administration and maintenance of Showcase 9 software. For information about installing the product, see the *Showcase 9 Installation Guide*.

About This Guide

- **Chapter 1, "Technical Support,"** contains contact information.
- **Chapter 2, "Starting and Stopping the Server,"** explains how to start and stop the server and how to prepare the server for client connections.
- **Chapter 3, "Controlling and Managing Performance,"** explains how to manage performance to ensure successful deployment.
- **Chapter 4, "Managing Data on the IBM i,"** describes how to maintain your server library, including backing up and restoring files and checking for usage.
- **Chapter 5, "Server Maintenance,"** describes typical server maintenance procedures.
- **Chapter 6, "Understanding Report Services,"** describes Report Services, which runs reports and publishes results without requiring Enterprise Server.
- **Appendix A, "Commands,"** contains a list of commands commonly used to maintain IBM i libraries.
- **Appendix B, "Sample Databases,"** describes the sample databases installed with the server.
- **Appendix C, "Troubleshooting,"** lists common problems and questions regarding Showcase 9 applications and provides resolutions and answers.

Additional Resources

- For information about installing and using IBM Collaboration and Deployment Services, see the documentation provided by IBM.
- For tasks and steps for Showcase features unrelated to the installation and maintenance of IBM i server software, see the online help in each application.
- For information on configuring TCP/IP, see the *Showcase 9 Installation Guide*.
- For information on upgrading to Showcase 9, see the *Showcase 9 Upgrade Guide*.

Table of Contents

| | |
|---|-----------|
| Technical Support | 1 |
| Contacting Technical Support | 1 |
| Technical Support Website | 1 |
| Starting and Stopping the Server | 3 |
| Overview | 3 |
| Starting the Showcase Warehouse Manager Server | 3 |
| Stopping the Showcase Warehouse Manager Server | 3 |
| Starting and Stopping the Enterprise Server | 4 |
| Controlling and Managing Performance | 5 |
| Managing a Relational Database | 5 |
| Planning Security | 5 |
| Controlling Showcase Query and Showcase Report Writer System Usage | 12 |
| Controlling Level of Parallelism | 15 |
| Enabling Expert Cache | 16 |
| Tuning System Performance | 16 |
| DB2 Symmetric Multiprocessing for IBM i | 17 |
| Using Auditing to Understand Your Queries | 17 |
| Improving Query Performance | 21 |
| Optimizing Performance for Showcase Warehouse Builder | 25 |
| Managing Data on the IBM i | 29 |
| Moving Data from a Test to a Production Environment | 29 |
| Saving and Restoring Data | 29 |
| Redirecting Data | 29 |
| Creating Custom Library Lists for Users | 30 |
| Restoring Data | 31 |
| Disaster Recovery Backups | 31 |
| Changing Physical Files That Have Dependent Data Views | 31 |
| Changing Data Views That Have Dependent Data Views | 32 |
| Managing IBM i Passwords | 32 |
| Changing Passwords with an Showcase Application | 32 |
| Updating Passwords with the SCCHKPWD Command | 32 |
| Managing Information Workflow | 33 |
| Running Distribution Sets with STRDD and SBMDD | 33 |
| Running Distribution Sets from within IBM Collaboration and Deployment Services | 34 |
| Multiple QAQQINI File Support | 35 |
| Showcase Integration with JD Edwards (JDE) | 36 |
| Non-IBM i Profile Sign-On | 36 |

| | |
|--|----|
| User-Defined Codes (UDC) | 36 |
| JDE Security | 36 |
| JDE Multiple Library List Environments | 42 |
| JDE User Authentication in IBM Collaboration and Deployment Services | 45 |

Server Maintenance 49

| | |
|--|----|
| Using TCP/IP | 49 |
| Configuring TCP/IP with STRSCSVR | 49 |
| Finding Available Port Numbers | 49 |
| Migrating Users to a New Server Library | 49 |
| Configuring Firewall Settings with ADDSCTCP | 50 |
| Checking Server Library Usage | 50 |
| Listing Connections to Relational Applications Only | 50 |
| Listing All Connections to All Applications | 51 |
| Ending TCP/IP (Stopping the Server) | 51 |
| Uninstalling a Server Library | 51 |
| Reloading the Server Configuration | 52 |
| Backing Up Server Information | 53 |
| Saving Server Information | 53 |
| Restoring Server Information | 54 |
| Restoring Server Information | 54 |
| Detecting and Removing Obsolete Information | 59 |
| Restricting Access to a Server Library | 59 |
| Applying the Restricted State | 59 |
| Removing the Restricted State | 60 |
| Unapplying Server Patches | 61 |
| Working with Server Logs | 61 |
| Working with Journals | 62 |
| Single Sign-On Configuration (SSO) | 62 |
| Enterprise Identity Mapping (EIM) Configuration | 62 |
| Showcase Warehouse Manager Server Configuration | 63 |
| Showcase IBM i Data Source Configuration | 64 |
| Showcase Warehouse Manager Server and IBM i Work Management | 64 |
| Work Management Controls | 64 |
| Security Issues | 67 |
| Command Authorities | 67 |
| User Exit Programs | 67 |
| Distributed Relational Database Architecture (DRDA) | 69 |
| Configure Showcase 9 Server to Reject Requests from Down-level Clients | 69 |

Understanding Report Services 71

| | |
|---|----|
| Report Services for IBM Collaboration and Deployment Services | 71 |
| Process | 71 |
| Requirements | 72 |
| Localization Settings | 72 |
| Image and Chart Support | 72 |

| | |
|--|------------|
| Data Sources | 74 |
| Default Format Template Definition | 74 |
| Overview | 74 |
| Configuration | 75 |
| XML Template Files | 75 |
| Sample Formats | 76 |
| Compatibility and Limitations | 79 |
| Appendix A - Commands | 83 |
| IBM i Commands | 83 |
| Showcase Warehouse Manager Server Commands | 83 |
| SCCHKPWD Command Parameters | 87 |
| Appendix B - Sample Databases | 89 |
| Installing the Sample Databases | 89 |
| Contents of the SCSAMPLE90 Sample Database | 89 |
| SCSAMPLE90 Database | 90 |
| Installing the Sample Library after Installation with RSTSMPLB | 97 |
| Appendix C - Troubleshooting | 99 |
| Installation Troubleshooting | 99 |
| Showcase Query Troubleshooting | 99 |
| Showcase Report Writer Troubleshooting | 101 |
| Showcase Warehouse Builder Troubleshooting | 101 |
| Showcase Warehouse Manager Server Troubleshooting | 101 |
| Index | 103 |

Technical Support

Contacting Technical Support

For general Showcase Information

Help/Systems can be reached by calling 952-933-0609.

For technical support or information

Call our general number 952-933-0609, and ask for technical support.

-or-

Send an Email to showcase@helpsystems.com.

For information on Showcase products, services, and partner programs

Visit the the Showcase home page:

www.helpsystems.com/showcase

To download documentation, software, or the latest program fixes

Go to the Showcase home page:

www.helpsystems.com/showcase

Technical Support Website

Customers who are current on maintenance can apply for a personalized ID for the Help/Systems Technical Support website. If you do not already have an ID, access the Technical Support website at www.helpsystems.com/showcase and apply for an ID. Once you log in to Technical Support, you can access information such as the following:

Troubleshooting. Search Help/FACTS for a resolution and view frequently asked questions.

Documentation. Read technical support policies, browse the technical support newsletters for technical issues, or view and download the latest product documentation.

Supported Platforms. Find supported software versions, release dates, end dates, IBM i and PC operating systems, and web browsers.

Upgrade Info. Order an upgrade and view upgrade information pertinent to each release.

Patches. Download patches for all supported products.

Contact Us. Find the most current technical support phone numbers, e-mail addresses, and fax numbers.

Other Links. Find information about products, training, and upcoming events.

Starting and Stopping the Server

Overview

This chapter outlines the procedures for starting and stopping Showcase Warehouse Manager Server and Enterprise Server.

***Note:** Stop the servers before performing an IBM i task that may jeopardize system stability (such as applying PTFs).*

Starting the Showcase Warehouse Manager Server

Running STRSCSVR starts the Showcase Warehouse Manager Server. The Warehouse Manager Server TCP/IP server job should start automatically when the installation is finished, provided that TCP/IP was configured correctly.

However, if you stopped the server, you must manually start it with the STRSCSVR command.

1. Log on with an IBM i profile that has *JOBCTL special authority.
2. Run the command SCSEVER/STRSCSVR, where SCSEVER is the library name in which the server is installed.

Stopping the Showcase Warehouse Manager Server

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

1. . Log on with an IBM i profile that has both *ALLOBJ and *JOBCTL special authorities.
2. Add the server library to your library list.
3. Run the command SCSEVER/ENDSCSVR, where SCSEVER is the library name in which the servers are installed.

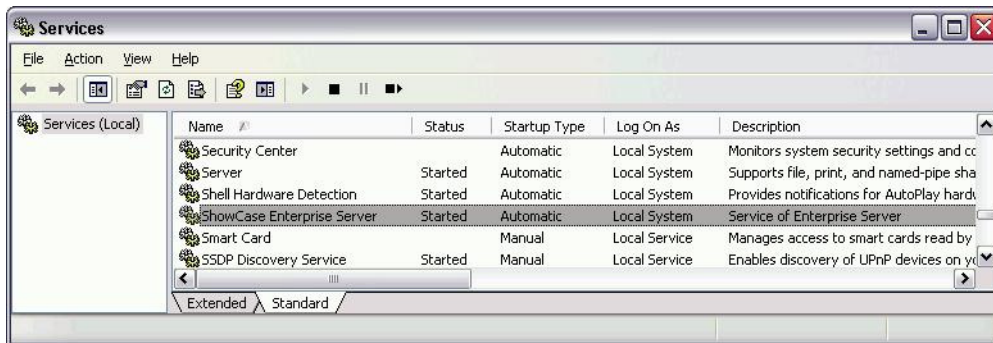
The ENDSCSVR command terminates all current TCP/IP connections to the specified library, but it does not release all locks on Showcase objects. Before performing Showcase server maintenance (applying a patch, for example), please refer to the maintenance instructions to end object locks.

For more information about uninstalling server libraries and related topics, see “Uninstalling a Server Library” on page 51 in Chapter 5.

Starting and Stopping the Enterprise Server

Showcase Enterprise Server runs as a Windows service program in Showcase 9, allowing the server to start automatically without user interaction. Previously, reporting jobs could be delayed indefinitely (for example, after a reboot) until a PC user with account access started the server manually. As a Windows service program, Showcase Enterprise Server is set (by default) to start automatically when the system starts. From the Properties dialog box, the startup type (Automatic, Manual, or Disabled) can be set to your preference, and the service can be manually stopped, restarted, paused, resumed, and disabled. To configure Enterprise Server service settings, choose **Start \ Control Panel \ Administrative Tools \ Services** and then double-click **ShowCase Enterprise Server**.

Figure 2-1
Running Enterprise Server as a Windows service



Enterprise Server can also be started by choosing **Start \ [All] Programs \ IBM \ ShowCase 9 \ ShowCase Enterprise Server**, but this method should not be necessary if the Windows service program starts automatically.

There are two logs available for Enterprise Server:

- Wrapper log. This log is enabled by default. The wrapper log file captures information displayed in a MS-DOS console window after executing `scserver.exe` or starting Enterprise Server through the menu option. The log is generated at `<install_dir>\logs\wrapper.log`.
- Debugging log. This log is disabled by default. To enable the log and capture debugging information, stop the Enterprise Server, set the following properties in the `showcase.cfg` file, and then restart the Enterprise Server. The log is generated at `<install_dir>\bin\scserver.log`.

```
showcasecorp.traceService=ALL
showcasecorp.traceLevel=ALL
showcasecorp.traceUser=ALL
showcasecorp.traceLog=scserver.log
```

Controlling and Managing Performance

Managing relational database performance can be a challenge for the most experienced database administrator. This chapter provides insight into performance planning and tuning relational databases.

These sections focus on performance tuning a relational database:

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

Refer to your IBM i documentation for the latest information on command parameters, optimizing performance, and so forth.

Managing a Relational Database

This section focuses on tuning relational database performance.

Planning Security

This section outlines the security issues to consider when deploying Showcase for relational data access.

IBM i Security Considerations

When deploying Showcase Warehouse Manager Server products, consider your current IBM i security environment. Items to address when creating a security plan include:

- What information do you currently store or plan to store on your system?
- Will you use Showcase applications on an IBM i system that is also currently used for transaction data processing, or will you use Showcase software on an IBM i system that is 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 nonrestrictive security approach? Do you use a firewall?
- 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 Server products?

- What other vendors' applications affect the way you set up your system security?
- What IBM i system security level will you use?

It is important to plan your overall IBM i security approach or have a thorough understanding of the existing IBM i security approach. This allows you to select the best Warehouse Manager Server security implementation to meet your current and future security needs. Refer to the IBM i documentation for more detailed system security information.

Server Security

When you set server security, it affects all client applications connecting to the server. These include:

- Showcase Warehouse Manager Client
- Showcase Warehouse Builder (Showcase application security only)
- Showcase Query
- Showcase Report Writer

Several factors affect server security. These factors interact with your current IBM i security settings, your current applications, applications purchased from other vendors, and your additional security measures (such as exit programs) to provide your total security solution. The primary Warehouse Manager Server security factors include:

- Using Warehouse Manager Client to apply library security, table security, column security, or row security for specific users or user groups.
- 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 IBM i objects.

The following sections make Warehouse Manager Server security plan recommendations based on specific situations. In general, you will always follow these steps:

1. Designate at least one user profile as a data administrator.
2. Review your user-profile scheme and update it if desired.
3. Select a general security policy. Choose whether to restrict global access to files and then authorize specific users to access specific files, or whether to allow global access to files and then exclude specific users from specific files. Implementing this model requires the following:
 - Set the server options to allow all users to see all libraries and tables or to exclude all users from all libraries and tables.
 - Use Warehouse Manager Client to apply authorities to specific tables for specific users.
4. Decide whether to restrict users to libraries within their IBM i library lists.
5. Apply any necessary column filters.
6. Apply any necessary row filters.
7. Ensure that the correct users have access to, or are restricted from, specific Showcase applications.
8. If desired, use Warehouse Manager Server products to save 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) are able to view all libraries and tables in Warehouse Manager Server products; they are also able to set server options. If you choose a data administrator user profile without *ALL-OBJ authority, you must also complete the following steps:

1. Using a profile with *ALLOBJ authority, use Warehouse Manager Client to authorize the user to Warehouse Manager Client software. See the Warehouse Manager Client online help for more information.
2. Using a profile with *ALLOBJ authority, give the data administrator IBM i *ALL authority to the user profiles that he or she will administer. For example, if the data administrator administers all user profiles, use the following IBM i command:

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

3. Using a profile with *ALLOBJ authority, give the data administrator IBM i *ALL authority to the libraries and tables that he or she will administer. For example, if the data administrator administers all libraries, use the following IBM i command:

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

4. Using a profile with *ALLOBJ authority, give the data administrator IBM i *USE authority to the following commands in the IBM i server library:
 - CHKSRVRCAT
 - CHGSRVRDFT
 - CHGSRVRSTS
 - SETLIBACC
 - DSPSCVER

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

```
CHGSRVRDFT :
```

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

5. Using a profile with *ALLOBJ authority, give the data administrator IBM i *CHANGE authority to the object SCVERSION in the server library on the IBM i. This allows the data administrator to change the security and *ALLOBJ server options. For example, use the following IBM i command:

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

Advantages of designating data administrators: Designating data administrators allows users to administer Warehouse Manager Server security without requiring *ALLOBJ authority. If the user in charge of implementing Warehouse Manager Server security does not have *ALLOBJ authority, the user should be a data administrator.

Disadvantages of designating data administrators: Data administrators cannot be locked out of any data in any products (that is Query or Warehouse Builder) that use Warehouse Manager Server as their server. In other words, data administrators can see any object for which they have IBM i authority.

Reviewing and Updating Your User Profile Scheme

In several of the recommended practices that follow, you are 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 can use group profiles. If the same security settings apply to several users, such as all members of a given department, you can create a group profile and make those user profiles members of the group. You can then administer Warehouse Manager Server security for the entire group rather than for each individual user.

Use Warehouse Manager Client to create group profiles and assign users to the group.

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, users do not have access from an IBM i display session. In this scheme your data files will not have object-level security applied; that is, you have not excluded specific users from specific files. Almost all client-server tools can bypass this security. You no longer have a menu interface protecting access to your files outside of the IBM i display session environment.

Menu-level security is considered a nonrestrictive security plan. Your goal when deploying Showcase should be to restrict users from gaining access to restricted data.

Option 1-To restrict all access and authorize users only to files they need:

1. 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
2. Using Warehouse Manager Client, authorize specific users or groups to the files they need.

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

Disadvantages. The process of authorizing users 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 by a user or group. 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. You would then need only to authorize a user to a library to allow access to all tables in the library.

Option 2-To restrict access to specific files:

Using Warehouse Manager Client, restrict specific users or groups from files 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 room for error and is not recommended. If you do not remember or do not 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 can read them until you restrict them.

Option 3-To run at security level 20:

Running an IBM i at security level 20 is a variation on menu-level security. At security level 20, all users have *ALLOBJ authority, meaning they can access all information in your data files. If this is

the case, you are probably relying on menu security to restrict file access. However, outside of your IBM i display session environments, users are not restricted. If this is the case, you need to take additional measures to restrict system access.

1. Using Warehouse Manager Client, change server options to Enforce Showcase Security on *ALLOBJ profiles.
2. Decide on the best approach for your situation. If you want to globally restrict access and authorize users to specific files, continue by following the steps in option 1. If you want to globally allow access and restrict users from specific files, refer to option 2.

Advantages. This is the only method for allowing Warehouse Manager Server security settings to take effect when the system security level is set to 20.

Caution: Menu-level security methods are effective when users are using Showcase applications. However, other ODBC products, iSeries Client Access, or FTP still allow users access to files and libraries. These applications bypass your security scheme because they do not tie into Warehouse Manager Server's security; there is no underlying object-level security.

Application-Level Security

With an application-level security scheme, all tables use object-level security that excludes user access. 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 not designed to adopt the correct authority, they are excluded from the tables.

Application-level security is considered a restrictive security plan. Your goal when deploying Showcase should be to allow users access to the tables they need but still restrict them from using tables they should not see.

Security on a New, Dedicated Data Warehousing IBM i

If you deploy security on a new IBM i that is dedicated to your data warehousing applications, you will not need to consider an existing security scheme. We recommend you use a restrictive approach when deploying a new data warehousing system.

- Run the IBM i at system security level 30 or higher.
- Do not grant users *ALLOBJ authority unless necessary for system operations or other non-Showcase 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 Server products.

Advantages. This approach allows you to control Warehouse Manager Server security without the risk of users accessing unauthorized database files. If you use other applications on the IBM i in addition to Showcase applications, you can set authorities for specific files needed by those applications without affecting Warehouse Manager Server security.

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

Restricting Users to Libraries in their IBM i Library Lists

Library list security allows you to restrict user access to the libraries in their library list. Without library list security, users have access to all libraries.

To set library list security:

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

Advantages. If you are currently using the job description library lists 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 JD Edwards package uses library-list security in some cases.

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

Applying Column Filters

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

To apply column security with Warehouse Manager Client:

1. Ensure that users have authority to see the database file and highlight the column you want to protect.
2. 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, select the appropriate user profiles and specify the Exclude authority setting.

Advantages. Without column security, you cannot restrict users from file data that contains columns with sensitive information. Column security gives you more flexibility in deciding which files your users access.

Disadvantages. Warehouse Manager Server column security applies only when using client products that use Warehouse Manager Server products.

Applying Row Security

Warehouse Manager Server row security allows users to view specific database row data while restricting them from viewing 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 Server row security, you can exclude users from seeing rows that contain information regarding prospects in other countries.

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

***Note:** The row security field is limited to 1024 characters. However, when Unicode characters are entered, the field limit is closer to 400 characters, due to the Unicode character conversion process.*

To apply row security with Warehouse Manager Client:

1. In our example, you ensure that each sales manager has access to the prospect file, highlight the name of the table, and choose row security.
2. 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. You could, for example, select 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. User exit programs specified on the PCSACC parameter can be used to apply row security for all client-server programs that support this parameter.

Advantages. Row security provides more flexibility in deciding which files users may access. Row security enables you to restrict user access to data in files that contain rows users should not see.

Warehouse Manager Client allows you to specify row security through a graphical user interface, without requiring user exit programs. It also allows you to apply row security for Showcase applications without affecting other applications.

Disadvantages. Warehouse Manager Server row security is effective only when you are using client products that use the server.

Setting Application Security

Warehouse Manager Client allows you to specify which users can run Showcase applications. Other 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 the applications.

Authority to use Query and Report Writer can generally be granted 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 Application Output

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

- Ensure that new objects do not unintentionally expose data to users who were previously restricted from viewing data
- Ensure that existing data are not accidentally overwritten

Ensure That Restricted Data Are Not Exposed

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

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

The same situation can occur with Warehouse Builder target files. It is recommended that you exclude users from access to these applications. Only those who deploy data warehousing or administer database access should use Warehouse Builder.

Ensure That Existing Data Are 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 prevents a user from accidentally replacing a production file 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 replacing the results.
- You must have system object-level security and at least Read authority to the target table before replacing 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, consider using the first mode to gain additional Warehouse Manager Server security. Users will not be able to accidentally replace production data with query results.
- If you have appropriate object-level security, you can disable the additional Warehouse Manager Server security.

To change batch output security:

1. Open Warehouse Manager Client, and from the Manage menu, choose Server Options.
2. In the Server Options dialog box, on the Batch tab, 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 IBM i*CHANGE authority to the table

Controlling Showcase Query and Showcase Report Writer System Usage

You can improve system performance by controlling how Showcase Query and Showcase Report Writer use system resources. This entails managing certain aspects of your IBM i to control the impact your PC applications have on the system. This is accomplished by setting client session timeouts, 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 Showcase Warehouse Manager Client.

Setting Showcase Query and Showcase Report Writer Session Timeouts

By default, timeout limits are not set. Administrators can optionally set session timeout values in Showcase Warehouse Manager Client so Showcase Query, Showcase Report Writer, and Excel/Lotus add-in users are automatically disconnected when their session is inactive for the specified time period (30-minute minimum). The timeout value can be set at a global, group, or individual user level. The timeout value applies to both Query and Report Writer, but if one of the applications has not exceeded the timeout value, then the timeout will not occur. When the timeout limit is reached, the client applications are disconnected from the data source(s), a timeout message appears, and users are prompted to save unsaved data. Administrators can set a different timeout value for each data source. When multiple data sources are connected, the smallest timeout value of the con-

nected data sources is used. The timeout setting does not apply to Showcase Warehouse Builder or Warehouse Manager Client sessions.

***Note:** When reports and queries run against the server, the client sessions are considered active until execution completes. Timeouts will not occur during this period of execution.*

To Set Client Session Timeout Limits

1. In Showcase Warehouse Manager Client, log into the server.
2. Right-click the applications tree node and choose Explore.
3. Right-click the Showcase Query application and choose Resource Settings.
4. Add a new entry by clicking the plus (+) icon. The Add Resource Settings for Profiles dialog box appears.
5. Select the profile for which the timeout value should be set (for example, select *PUBLIC for all users and groups).
6. Click Next.
7. Select the Use defaults radio button to accept the default settings.
8. Click Finish. The new entry will be added to the Resource Settings window.
9. Right-click the newly added entry and choose Edit.
10. Set a customized timeout value. By default, the timeout value is set to infinity (no timeout).
11. Click OK.
12. Click Update to save the settings.

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 IBM i predictive query governor support. Using Showcase Warehouse Manager Server, you can set the maximum run length for each user's query. If the query optimizer on the IBM i predicts that a query will take longer than the specified limit, it stops the query before it is run.

Advantages. Setting query run limits allows you to control the system resources allocated to query users. 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 query runtime, it may not be completely accurate (the query could actually have a shorter or longer runtime than the estimate).

Assigning Query Run Limits

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

1. Determine the average amount of time it takes to run several typical queries. You can use the Performance Analyzer feature in Showcase Query to estimate the amount of time a query will take to complete, or you can use auditing in Showcase Warehouse Manager Client and run the Detailed Select Statement query.
2. Determine the other system requirements and the job run priorities established earlier.

3. Using this information, break your queries into two groups:
 - Queries that complete quickly enough to be run interactively
 - Queries that should run in batch
4. Use Warehouse Manager Client to set up the appropriate query limits for *PUBLIC, individual users, and user groups.

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 QUSER:

- For queries estimated to run less than 5 minutes, run them interactively at priority 50.
- For queries estimated to run more than 5 minutes, run them in batch at priority 50.
- For queries estimated to run more than 30 minutes, do not allow them 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 are allowed to run. Also, some queries may not be allowed to run even though they finish within the specified time. To verify whether a query runs within the limits, use the Performance Analyzer in Showcase Query, which informs you of the query's estimated runtime.

Setting Server Job-Run Priorities

The job-run priority is used to prioritize the server job relative to other system processes. The job-run priority is the standard IBM i 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 system processes.

Disadvantages. Another process can delay the query if it has a higher priority.

Assigning Job-Run Priorities

1. Consider all processes on your system and rank them by importance.
2. Based on ranking, choose the job-run priority you want to use when running queries.
3. Using Showcase Warehouse Manager Client, assign job-run priority for *PUBLIC.

You may further refine priorities by assigning job-run priorities based on user ID or user groups. This allows you to prioritize query usage among users.

Example

After evaluating your system, run the following jobs:

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

Assign the following job-run priorities:

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

The order-entry application and batch processing application have predefined run priorities. For queries, define a range allowing further prioritization of individual users as needed.

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

Additional Information

Once the server is started, the job priority can be changed using the CHGJOB CL command. When a user connects to the IBM i or submits a query to batch, the following attributes determine which run priority is used:

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

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

For batch jobs, the Warehouse Manager Client 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 feature on your IBM i. Controlling the level of parallelism determines whether the system allows a query to be logically split into multiple, independent tasks that can be carried out in parallel.

The IBM i system value QQRYPDEGREE can be used to tune the level of parallelism when running queries on your system.

Assigning the Level of Parallelism

1. Consider all work performed on the system and decide if queries should be allowed to consume additional CPU and I/O resources in order to finish quickly.
2. Decide on an appropriate amount of parallelism to allow. Currently, the system value provides the following control:
 - *NONE. Do not allow any parallelism. Queries are not the highest priority and the system will not devote the necessary resources to an individual query.
 - *IO. Allow I/O parallelism. The system uses additional I/O and mainstore resources to complete a query faster. In some cases, allowing the system this freedom to use more resources while running a query improves the cumulative total resources used.

- *OPTIMIZE. Allow the system to choose whether I/O 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 improves the cumulative total resources used.
 - *MAX. Allow I/O and CPU parallelism. Parallelism is used whenever possible if it will help to improve the individual query. The system is used only for running queries, so there is little impact on other applications.
3. Once the level of parallelism is established, review the system impact and adjust until you have the setting that works best for your environment.

Example

Your company has three systems. Queries are run on all systems, so you purchased DB2 Symmetric Multiprocessing for IBM i to benefit from the additional data access methods and you now want to use the appropriate level of parallelism. The following shows how each company system is used:

- SYS-A. This system runs an order-entry application and 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 the queries run.

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

- SYS-A=*NONE. The order-entry application is the 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. The system decides. 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, query completion time is improved and system paging and faulting rates are reduced. Enabling expert cache should be the first step in tuning your system for running queries.

To Enable Expert Cache

1. Determine the storage pool the server job is using. Normally, this is *BASE (storage pool 2).
2. On the IBM i, use the WRKSYSSTS command with the assistance level set to 2=Intermediate.
3. 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, the following is recommended:

- Always consider the entire system and all system work when adjusting systemwide 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 IBM documentation related to work management and performance tools, or an IBM i consultant specializing in system performance, to help get you started.

DB2 Symmetric Multiprocessing for IBM i

The DB2 Symmetric Multiprocessing for IBM i feature improves query performance by adding the following to IBM i's DB2:

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

The server takes advantage of the functionality provided by this feature. While the name indicates that it is for multiprocessor systems, it also provides improvements on single-processor machines.

Using Auditing to Understand Your Queries

Showcase can provide you with information about the impact users have on a system when accessing the IBM i relational databases using SQL. To collect this information, you must use Showcase Warehouse Manager Client and choose Auditing from the Manage menu. This information can help you make tuning and performance decisions regarding specific queries.

Impact to Overall Performance

Turning on the auditing feature impacts a system's performance. 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 system performance. Because environments differ greatly, it is difficult to predict the exact system performance degradation.

Because of the possibility of performance degradation, it is recommended that auditing be used only for a limited time. Collect information for a period of time, turn auditing off, and analyze the collected data. Do not leave the auditing feature permanently turned on. Auditing should be used as part of a planned cycle of information collection and analysis.

Collecting Information

1. Using Showcase Warehouse Manager Client, connect to the IBM i data source (the server) in which you would like to audit queries.
2. From the Manage menu, choose Auditing.
3. 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 is collected on SQL SELECT statements for each user who connects to the server library. These will probably be Showcase Query and Showcase Report Writer users. Local distributions run by Showcase Warehouse Builder are also recorded.

Note that some applications, 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 applications have on your SQL usage, these statements are also audited.

In some cases, the server is able to optimize a user's Query or Report Writer query so IBM i 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 collected information, it is recommended that you begin analysis only after auditing has ended. Otherwise, your own queries on the auditing information are added to the database file as you analyze it. This could skew your view of your users' actual SQL usage.

In Showcase Warehouse Manager Client, you will find several queries designed to help you analyze audit data. From the Reports menu, choose Audit Queries. Each query is provided in both Showcase Query (.dbq) format and Showcase Report Writer (.rpt) format. An explanation of each query follows, with some suggestions on how to use them.

SELECT Statements: Basic Information

This report returns basic auditing information about the queries run. 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 audit information about the queries run. Use it when you want to start analysis of the SQL SELECT statements that have run on the system. You may choose to do further analysis after seeing detailed information on the Showcase applications used, the time to run statements, the query optimizer's access methods, and so on.

For each query run, there is at least one record in the report per file used in the query. Each record shows the access method the DB2 optimizer chose to use when retrieving data from the file. Always note the job number and unique statement number of the query you are interested in so that you can find information about that particular query in the other reports.

Ordering. For different analysis purposes, you can order the data in this report differently than the way it is presented. To order by different columns in the group, from the Format menu, choose and drag the column you want to order to the top of the Break Groups list.

Timings. The detailed information report includes the runtime for each query. The time reported is not measured in standard time; it is an internal processor time as reported by DB2. This information is best used for comparisons between queries.

This time may not exactly match 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 runtimes. The reason for this could 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 could be that the SELECT statement was run multiple times in the same job, and the application 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. 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 because the file might be joined to itself; it was actually used more than once. The query could 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. 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 when using Report Writer (.rpt) rather than Query (.dbq). The .rpt allows you to see multiple lines of the statement text.

Queries That Build Temporary Indexes

Use this report to find which queries the DB2 optimizer chose to build temporary indexes for and whether the optimizer advises building a permanent index.

Queries That Used Existing Indexes

This report returns a list of all SQL queries that used a preexisting system index. 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.

For some queries, 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 optimizer will occasionally advise 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 can improve 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 determine which are the most frequently queried files. This can assist you in making information more readily available to users or in deciding how to optimize your environment.

Which Queries Use a Given File

Use this report to find out exactly which SELECT statements ran against a file and which users ran them. To use this report, you need to enter the name of the file you are interested in and its library.

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

Query Governor Information

DB2 predictive query governor prevents queries from running if it estimates they will take longer than a specified time interval. 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.

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 predictive query governor when the specified user connects to the server. This report can help determine whether you are preventing too many queries from running due to the query limit being set too low.

Additional Analysis Tips

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

Managing Your Auditing Environment

Saving Auditing Results

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

1. Use the IBM i command CRTDUPOBJ to create a copy of the file SCDBMON, located in the server library in which auditing started. Specify *YES for the Duplicate data option.
2. Following the guidelines in the previous step, make a copy of the SCAUDINF file, also located in the server library.
3. Make a note of the new filenames and their respective libraries.
4. Open the auditing query or report that will run against the copied data.
5. From the File menu in Showcase Query or Showcase Report Writer, choose Save As and save a copy of the query or report.
6. Using the new, saved copy of the query or report, choose Tables from the Query menu. For each reference to either SCDBMON or SCAUDINF, click the table name, and then click the Change Link button. In the Change Link dialog box, enter the library name and the filename for the copy of SCDBMON or SCAUDINF you made earlier. You may leave any references to SCAUDMAP alone, because the file is neither cleared when you begin auditing nor updated when users run queries.

7. Repeat the previous step for each query or report that will run against the saved audit files.
8. Remember to delete your copies of SCDBMON and SCAUDINF when finished. The copies you save are no longer under the application's control, therefore the server will not automatically clean them up.

Subsequent Installations

If you reinstall the server in a library where auditing has run (for example, to upgrade the library to the next release of the server), be aware that your auditing information is not preserved during the installation. If you want to save your data, follow the steps in the previous section, being sure to save the files in a library other than your server library.

Auditing File Size

The auditing file SCDBMON can become quite large. 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 interval expires, be aware that any users currently connected when these parameters are reached continue to have their queries audited and information recorded in the auditing file until they disconnect.

Upgrading IBM i

Auditing should always end before you upgrade your IBM i version. If you want to preserve any previous auditing information, follow the steps at the beginning of *Managing Your Auditing Environment* on page 20. 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 Showcase Query and Showcase Report Writer system usage, it is possible to fine-tune your queries by taking advantage of methods used by the IBM i to process queries and by following some general query creation tips.

Understanding the Query Optimizer

The IBM i uses the query optimizer to efficiently retrieve end user and application data. Understanding how the query optimizer processes queries helps users create queries that 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 IBM i. 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 runtime code for data retrieval
- Passes those structures to the database runtime 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 and determining how the database functions access and process each record in the database.

A key piece of information held by the access plan is the query's access path. The access path refers to the method(s) the database management system (DBMS) uses to perform the query.

During the optimization phase, the query optimizer evaluates 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 most efficient method for processing the query.

For a more detailed explanation of how the query optimizer processes information, see the IBM documentation related to DB2 for database programming.

Using Showcase Query's Performance Analyzer to Evaluate Your Queries

Improving query performance is often 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 performance level.

To tune your query:

1. Understand how the query optimizer chose to implement the query.
2. Analyze the query runtime performance and determine if it is I/O or CPU bound.
3. 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.
4. Go back to step 1 until the performance does not improve or all possibilities have been attempted.

The most important factor 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.

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

To analyze your query:

1. After creating a query, from the Run menu, click Performance. The Performance dialog box displays.
2. Click Analyze to analyze the query's efficiency.
3. Once the analysis is complete, click the Log tab to view the query's implementation messages. Definitions for each message can be found in the Performance Analyzer online help.

Refer to the following sources for details regarding log information:

- Online help, which contains a description of the provided information along with some general performance-tuning tips
- IBM documentation related to DB2 for database programming
- IBM documentation related to DB2 for SQL programming

Query's Performance Analyzer allows you to quickly and easily test queries before they run to determine how they were implemented and the estimated runtime. This estimate is provided by the DB2 query optimizer and is fairly accurate. However, as mentioned previously, the optimizer does not always have adequate information for an accurate estimate. Note that Performance Analyzer is also available in Showcase 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 is updated for every insert, update, and delete operation. Indexes built over large files will also consume additional DASD.

Deciding on when an index should be built 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). This section focuses 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 query performance:

1. Look at a handful 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.
2. Identify the following query attributes:
 - Join conditions
 - Group-by columns
 - WHERE clauses
3. If the query messages gathered in step 1 point to join conditions or group-by clauses as the source of spontaneous index builds, consider building permanent indexes above those files. If an index does not 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.
4. 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.

Refer to IBM documentation for more information on creating indexes.

The following examples present sample SQL statements and the index types that could 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 depends on which is most selective (put the most selective first).

Example 3

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

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 builds an index at runtime if one does not exist over custname. This is fast when custnum is very selective because the existing index for custnum contains few records. The index for custnum is used to build the index for the group-by clause.

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 tables. If no relationship is defined, the system performs 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. If possible, it is a good idea to break the statement down into simpler queries that run back-to-back.

Limit Use of Data Type Conversion

Whenever possible, avoid forcing the system to perform data-type conversions. For example, if you have a WHERE clause, such as WHERE INTEGER = 1.4, the system must convert the INTEGER field into another data type to accurately compare it to 1.4.

Avoid String Padding

In cases where a string literal is compared to a field, ensure that 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 perform an accurate comparison. As with numeric conversion, this eliminates use of any index built over the character field. A shorter literal string does not cause problems, since the system pads the literal appropriately to use the index.

Minimize the Use of LIKE

When possible, avoid using pattern matching. 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 index use that would otherwise apply to the selection criteria.

Optimizing Performance for Showcase Warehouse Builder

Following are several ways to improve Showcase Warehouse Builder relational database distribution performance.

Choose the Most Efficient Distribution Type

Showcase Warehouse Builder provides several types of distributions. The various distribution types are shown in the following figure. The distribution type you choose has a significant impact on the performance of your distribution and its affect on your IBM i system. The performance characteristics of the various distribution types are described in Table 3-1.

Figure 3-1

Showcase Warehouse Builder distribution types




Table 3-1

Distribution types and performance levels

| Distribution type | Performance level |
|--|-------------------|
| Local (within the same machine) | Very good |
| Pull (from a remote IBM i data source to a target table on the control server) | Very good |
| Push (from the control server to a remote IBM i) | Not very good |
| Pass through (from a remote IBM i to another remote IBM i) | Not very good |

Use Pull Distributions

The fastest distribution type 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. It is recommended that all performance-critical distributions pull data from the source system. This means the control server and the destination server should be on the same system.

Creating Libraries

Showcase Warehouse Builder creates libraries on the target system by default to store distributed data. This improves distribution performance. To use journaling, you can specify that Warehouse Builder create collections instead. However, using collections and activating journaling could dramatically reduce distribution performance.

To create collections or libraries with Warehouse Builder:

1. From the Tools menu, choose Options.
2. On the Default tab, select Create Target SQL Collection. If this option is not selected, Warehouse Builder may create libraries.

Using CLRPFM Command Instead of Mass Deletes

When the target table is defined in Showcase Warehouse Builder, you can specify that the 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 works against any SQL data source. The SQL DELETE method is slower than CLRPFM but is applicable with all targets.

On the IBM i, 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 valid only with IBM i 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:

1. From the Tools menu, choose Options.
2. On the Default tab, in the Table Overwrite section, select the SQL DELETE option or the Clear File Member Command option.

Sending Only Changed Records

Another method to improving distribution performance is to send only changed records between the source and target servers. While there is currently no automatic method for sending 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.

Using CURDATE()

If the file(s) referenced on a source system include dates, 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, you may be able to use this field to select only the recently inserted records. To use this method:

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

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

Other Performance Tuning Issues and Solutions

The efficiency of the `SELECT` statement running on the source system has 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 Showcase Query.

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

or

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

Based on this information, you can 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 Data on the IBM i

The server supports Showcase PC applications. Data produced by Showcase applications are stored on the IBM i. Administrators can 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 server maintenance.

Moving Data from a Test to a Production Environment

Showcase Query, Showcase Report Writer, and Showcase Warehouse Builder allow users to develop queries, reports, data views, and data distributions directly against IBM i or select non-DB2 databases. The environment in which these files are developed is critical to the success of daily database operations. Developing these files in a production environment can adversely affect the performance of other operations. To prevent this problem, develop files against test data, either on a separate IBM i machine or in a separate library. Once testing is complete, place your queries, reports, data views, and/or data distributions into production.

There are various methods for moving 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 IBM i. 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 “Restoring Data” on page 31 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 is necessary to modify the database structure by renaming libraries or tables or by removing or renaming columns. However, it is recommended that you keep structural modifications to a minimum. If you are performing many structural modifications, you should perform 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 PC applications determine that libraries, tables, or

columns are missing or if column data types have changed. Consider the following two scenarios to help decide how to structure an IBM i database:

Scenario 1: Testing and Production on Separate IBM i Servers

If testing and production are performed separate IBM i servers, 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 | FY07 | CUSTOMERS | CUSTNAME |
| | SYS1 | FY07 | ORDERS | SHIPDATE |
| To | SYS2 | FY07 | CUSTOMERS | CUSTNAME |
| | SYS2 | FY07 | ORDERS | SHIPDATE |

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

Scenario 2: Testing and Production on the Same IBM i

If testing and production are performed on the same IBM i, 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 are different.

| | System | Library | Table | Column |
|------|--------|---------|-----------|----------|
| From | SYS1 | SAMPLE | CUSTOMERS | CUSTNAME |
| | SYS1 | SAMPLE | ORDERS | SHIPDATE |
| To | SYS1 | FY07 | CUSTOMERS | CUSTNAME |
| | SYS1 | FY07 | ORDERS | SHIPDATE |

After testing is complete, use the Change Link feature to convert a query, report, data view, or data distribution to run against the production library. This feature is available in Showcase 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. Developers or testers can run against a library list that points to test libraries; users can run against a library list that points to production libraries.

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

1. Open Showcase Warehouse Manager Client and from the Manage menu, choose Server Options. The Server Options dialog box displays.
2. On the Library Access tab, select Return only libraries in user portion of the library list.

This simplifies the database for users because the data they need are 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.

Restoring Data

The following situations require you to temporarily back up server information:

- You need to install a new release in the server library.
- You need to uninstall (see an explanation of the UNINSTSC command in the section “Uninstalling a Server Library” on page 51) the server library but want to save the library's contents, such as security settings, queries, and data views.
- You need to move the server library from a test system to a production system.
- You need to move 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 IBM i.

For any of these situations, use the standard Save/Restore backup 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 Chapter 5 for details.

Using this procedure, you save server information to a file. If you are moving data between systems, you can send files electronically, which is convenient for remote sites.

***Note:** It is recommended you perform the regular backup procedure using the SAVSRVRINF and RSTSRVRINF commands (see Chapter 5 for details). However, for disaster recovery purposes, you can restore from a tape or other media device to the same machine using the following procedure.*

Disaster Recovery Backups

To back up server information for disaster recovery purposes, you should follow the same procedure as for other IBM i backups. Whether you use the IBM i 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, server information is backed up when you perform the system save. If you typically just back up individual libraries, you must also back up the IFS of the Showcase Warehouse Manager Server library. For complete details regarding backing up objects from the IBM i, refer to your IBM iBackup and Recovery documentation.

If disaster should strike and you are forced to restore the Warehouse Manager Server library and the associated IFS directory, you should run the UPDSRVPLIB command as follows:

```
<SCSERVER>/UPDSRVPLIB SERVERLIB(<SCSERVER>) OPTION(*FIX)
```

where <SCSERVER> is the name of your Warehouse Manager Server library.

Changing Physical Files That Have Dependent Data Views

Showcase Query lets you create and manage simplified views of any IBM i database. This simplified view, or data view, is saved on the IBM i as an SQL view. Since SQL views are saved on the IBM i 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 IBM i table. You can run these predefined queries and reports from the Tools menu in Query.

To change a physical file that has dependent data views, 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 page 29. You will want to run the `SAVSRVRINF` command on the IBM i 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 performed within Showcase 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 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 IBM i Passwords

Showcase Warehouse Builder distributions rely on consistent passwords to connect to the IBM i. When a profile tries to connect to an IBM i through an Showcase application, the password is passed through to log on to the IBM i. In time, changed and/or expired passwords can cause the Showcase control tables to contain incorrect user passwords. This can cause Warehouse Builder distributions to fail.

Changing Passwords with an Showcase Application

If the user's IBM i password has changed or expired, the user can update the applicable product 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 Showcase Warehouse Builder distributions.

Note: Most Showcase applications support 128-character, mixed-case passwords. JD Edwards One-World supports 10-character, uppercase passwords. For complete information on IBM i password guidelines and restrictions, see your IBM i documentation.

Updating Passwords with the `SCCHKPWD` Command

Use the `SCCHKPWD` command to update an IBM i user profile's password in the product control tables. 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 generate a list of the items matching the search criteria.
- Limit the generated list to matching entries for DB2 databases, Showcase Warehouse Manager Servers, and non-Showcase data sources.

Examples

- Update passwords for an IBM i RDB (Showcase Warehouse Builder distributions):


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

- Update passwords for a non-Showcase data source (scheduled reports and Warehouse Builder Oracle or SQL Server pulls):

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

Managing Information Workflow

Running Distribution Sets with STRDD and SBMDD

The STRDD (start data distribution) and SBMDD (submit data distribution) commands run Showcase 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 of 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 watch 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. An RSC6336 message is a status message and the program is not required to watch 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
```

Running Distribution Sets from within IBM Collaboration and Deployment Services

IBM Collaboration and Deployment Services users can control individual Showcase Warehouse Builder definitions within job steps without breaking up their existing Warehouse Builder sets. Administrator or analysts can select individual Warehouse Builder definitions within a Warehouse Builder set at the job-step level and use existing IBM Collaboration and Deployment Services functionality to conditionally control the flow, set notifications, and so on.

***Note:** Note: For details about creating and working with jobs, see the IBM Collaboration and Deployment Services Deployment Manager on-line help. For details about creating Warehouse Builder definitions and sets, see the online help within the Warehouse Builder application.*

Adding a Showcase Warehouse Builder Job Step

To add a Warehouse Builder job step in Deployment Manager:

1. From the jobs palette, select Warehouse Builder.
2. Click anywhere in the job canvas. The Warehouse Builder work step is added to your job.

Selecting a Set and Individual Definitions to Run

To define the Warehouse Builder job step, click the step in an open job. The Warehouse Builder Set tab contains information about the Warehouse Builder set to be associated with the job step. Use the drop-down list to select a Warehouse Builder set to associate with the job step. The list contains all Warehouse Builder sets currently available on the control server. After selecting a set, the corresponding definitions contained in the set display in the main table on the tab.

Select All (or None) to select (or deselect) all definitions to run. Use the expand (+) and collapse (-) icons to show or hide the individual definitions. Select (or deselect) the individual Warehouse Builder definitions that you want to be included in (or excluded from) the job step.

Defining Properties

Before the step can successfully run, properties of the Warehouse Builder job step must be defined.

To define general properties for a Warehouse Builder job step, click the step in an open job. Click the General tab to view or edit the step properties. Type a name for the step. The default name appends a *_step* suffix to the file name.

Important Notes

- Running Warehouse Builder jobs cannot be canceled via Deployment Manager. To cancel a running Warehouse Builder job, you must manually cancel the job via the IBM i command line.

- Warehouse Builder supports the use of an IBM i Independent Auxiliary Storage Pool (IASP) rather than the local IBM i database. If you plan to use Warehouse Builder with IASP, you must configure separate Showcase IBM i Driver data sources via the Microsoft ODBC Administrator: one for the Warehouse Builder Control Server (local database), one for the Warehouse Builder Source Server (local database or IASP), and one for the Warehouse Builder Target Server (local database or IASP). For more information, see the help on the Connection tab when creating Showcase IBM i Driver data sources.

Multiple QAQQINI File Support

Showcase supports multiple QAQQINI files. QAQQINI files are used to specify DB2 SQL options for queries running through the STRSQL command (refer to your IBM i documentation for details on the options that may be set in the QAQQINI files). By default, the QAQQINI file in QUSRSYS is used for the Showcase user in Showcase Query and Showcase Report Writer.

The Showcase administrator can specify an alternative QAQQINI for a user, a group of users, or the special *PUBLIC group.

To specify which default QAQQINI file is used for a user or a group:

1. Sign on to the appropriate IBM i.
2. From the command line, add the appropriate Showcase Warehouse Manager Server library to your library list using the ADDLIB command. For example, ADDLIB LIB (SCSERVER) , where SCSERVER is the server library name.
3. Enter the command CHGUSRQRYA and press F4 to prompt for parameter values.
4. For the first parameter, specify the appropriate user profile, group, or the *PUBLIC special group.
5. For the second parameter, type the library name where the alternative QAQQINI file is located. Note that the library must exist and there must be an existing QAQQINI file located in the library prior to running the CHGUSRQRYA command.
6. Press Enter to commit the changes. A message in the job log indicates whether or not the change was successful.

To remove a default QAQQINI file used for a user or a group, follow the instructions for specifying a default QAQQINI file as documented above, but type the special value *REMOVE for the QAQQINI library name.

The following items are usage considerations related to multiple QAQQINI file support and the query time limit for interactive and batch queries:

- In Showcase Warehouse Manager Client, the administrator can set the interactive and batch time limit for users or user groups. The *INIFILE option is available, allowing a QAQQINI file to specify the query time limit SQL setting. With the *INIFILE option, the default QAQQINI file for a user or a user group controls the query time limit.
- The order to obtain the interactive or batch query time limit for a user is as follows:
 1. The Warehouse Manager Client numeric setting is used if found.
 2. If *INIFILE is specified, the user's default QAQQINI file is used to retrieve the query limit.
 3. If neither of the previous values is present, the query limit defaults to the value retrieved from the system value QQRYTIMLMT.

If a query time limit is used for interactive or batch reports, DB2 enforces the time limit; Showcase does not control how DB2 enforces the query time limit.

Note: Refer to the IBM i documentation for more detailed information.

Showcase Integration with JD Edwards (JDE)

Showcase 9 supports JJD Edwards World A9.2, World A9.3, and EnterpriseOne ERP 9.0, 9.1, and 9.2. The following sections describe enhancements to product integration with JDE. For additional JDE information, such as data dictionary information, visit www.helpsystems.com/showcase and log on to the Showcase products support site.

Important: If your IBM i and JD Edwards user profile names are the same, the passwords for both must be the same as well.

Non-IBM i Profile Sign-On

Non-IBM i profiles cannot be used with Showcase Warehouse Builder or Showcase Warehouse Manager Client, and they cannot be administrated with Warehouse Manager Client. With the installation of the IBM Collaboration and Deployment Services - Server Adapter for Showcase, a security provider is installed to your IBM Collaboration and Deployment Services server. This security provider may be configured to allow JDE application users to log in and use the IBM Collaboration and Deployment Services environment. For more information, see the topic JDE User Authentication in IBM Collaboration and Deployment Services on page 45.

Previously, JDE OneWorld users needed an IBM i user profile for product applications to recognize their environment information. Administrators can now use the `SETJDEOWA` command to enable product applications to recognize environment information for OneWorld users who do not have an IBM i user profile.

Refer to the `SETJDEOWA` Command section on page 40 for detailed command information.

For JDE EnterpriseOne 9.2 configuration, see page 46.

Note: Most Showcase applications support 128-character, mixed-case passwords. JDE OneWorld supports 10-character, uppercase passwords. For complete information on IBM i password guidelines and restrictions, see the appropriate IBM i documentation.

User-Defined Codes (UDC)

During the initial JDE system setup, 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 columns display. UDC description fields can be used wherever columns are used. UDC customized descriptions are available in Showcase Warehouse Builder, Showcase Query, Showcase Report Writer, and when running interactive reports via a browser in IBM Collaboration and Deployment Services Deployment Portal.

When the UDC code field (not the description field) is present for a prompt, all possible UDC prompt values and descriptions are returned. When a prompt includes both a UDC code field and description field, only the distinct table data is returned for the prompt values.

To view the JDE UDC information in Warehouse Builder, the IBM i and JDE OneWorld profiles must have `BV3C` defined as the Initial Program to Call. `BV3C` is located in the IBM i directory `QGPL` and must be qualified.

JDE Security

The security rules for the specific JDE applications are as follows:

- For World and OneWorld, the security rules follow User, then Group, then *PUBLIC.
- For EnterpriseOne, the security rules follow User, then Role, then *PUBLIC.

The following sections discuss security options for the various JDE applications.

Application Security

JD Edwards OneWorld and EnterpriseOne provide options for implementing application security as table security. Application security can be applied to a set of tables or a super-set based on system code. Showcase honors this level of security in Showcase Query, Showcase Report Writer, and IBM Collaboration and Deployment Services Deployment Portal.

If a OneWorld/EnterpriseOne table belongs to multiple OneWorld/EnterpriseOne applications, the table is available if the user has run authority for at least one of the applications that use the table. Conversely, if one application is secured from the user, the table should not be available to the user. An option is required to support both options for tables belonging to multiple applications.

Application security for OneWorld/EnterpriseOne tables belonging to multiple OneWorld/EnterpriseOne applications is handled in the following manner:

- If the Run Authority Override parameter is set to *YES on the SETJDEOWA command, a table is available to the user if the user has run authority for at least one of the applications that use the table.
- If the Run Authority Override parameter is set to *NO on the SETJDEOWA command, the table is not available to the user if the user is restricted from running at least one of the applications that use the table.

The RUNAUTHOVR parameter on the SETJDEOWA command provides application security configuration options.

***Note:** Refer to "SETJDEOWA Command" on page 40 for more details regarding SETJDEOWA command parameters.*

Object-Level Security

Through the JD Edwards OneWorld client interface, via the Object Management Workbench, users and groups can be secured from database tables and table columns. The tables are secured by marking them as **available** or **not available**. Showcase honors this level of security in Showcase Query, Showcase Report Writer, and IBM Collaboration and Deployment Services Deployment Portal (running queries through a web browser interface).

In Query and Report Writer, the table list does not contain tables or columns that the user, their group, or the *PUBLIC group is not authorized to view. In addition, it is not possible to use an unauthorized table/column when this restriction is explicitly entered in the table selection dialog box for Query and Report Writer. Dynamic Queries check the table/column authority at runtime and prevent the current user from running the query if the current user, group, or *PUBLIC has been excluded from the table/column.

Column-Level Application Security

Column-level application security behaves in the same manner as JD Edwards application security, except security is honored at the column level. Showcase honors column-level application security-

as defined in OneWorld/EnterpriseOne-in Showcase Query, Showcase Report Writer, and IBM Collaboration and Deployment Services Deployment Portal.

When security is defined for columns, referenced within one or multiple tables, the security is honored for all tables that reference the restricted columns (provided that all tables are associated with the same OneWorld/EnterpriseOne application). Column-level security is automatically enabled when the SETJDEOWA command parameter Use JDE OneWorld App Security is set to *ENABLE.

***Note:** Refer to "SETJDEOWA Command" on page 40 for more details regarding SETJDEOWA command parameters.*

Column-Level Security

Column-level security functions in the same manner as column-level application security with the exception that the security rules can only be applied to columns referenced in a single table, as opposed to multiple tables.

Row-Level Security

OneWorld supports row-level security on any column. EnterpriseOne supports inclusive and exclusive row-level security on any column.

Inclusive Row Security

Inclusive row security allows the user to define valid value ranges in the F00950 Security Workbench Table. This ensures that users/groups in the security records only have access to records within the specified range and is the opposite of using exclusive row security where the defined ranges are those that the users/groups should not have access to. Depending on the security setup, inclusive row security can increase performance over exclusive row security. Performance can be improved when the use of inclusive row security results in a small range of valid values in the row security application (P00950), rather than specifying a large range of secured values in the row security application using exclusive row security.

At a system level, the decision must be made to use either inclusive or exclusive row security. The two security methods cannot be used in combination.

Activating Inclusive Row Security

Refer to the appropriate JD Edwards documentation for information on activating inclusive row security.

Scenario: Inclusive Row Security

The following identifies the required records that are created in Row Security for an example with inclusive row security. In this example, user MG5700778 will be able to view records in the Address Book table (F0101) with a business unit value between 1 and 20 and between 51 and 70. However, that user will only be allowed to update records in the Address Book table with a business unit value between 1 and 20. For inclusive row security, the following two records are added:

| User | Table | Data item | From Value | Thru Value | Add | Change | Delete | View | Alias |
|-----------|-------|------------|------------|------------|-----|--------|--------|------|-------|
| MG5700778 | F0101 | CostCenter | 1 | 20 | N | Y | N | Y | MCU |
| MG5700778 | F0101 | CostCenter | 51 | 70 | N | N | N | Y | MCU |

Selects performed against the F0101 table would resemble the following:

```
SELECT * FROM TESTDTA.F0101 WHERE ( ( ABMCU BETWEEN ' 1' AND ' 20' OR ABMCU BETWEEN ' 51' AND ' 70' ) ) ORDER BY ABAN8 ASC
```

Updates on the F0101 table would resemble the following:

```
UPDATE TESTDTA.F0101 SET ABALKY=' ',ABTAX='546',ABALPH='John Doe', ABDC='JOHNDOE',ABMCU=' 60',ABSIC=' ',ABUSER='MG5700778', ABPID='EP01012',ABUPMJ=101214,ABJOBN='DEN123456',ABUPMT=154030.000000 WHERE ( ABAN8 = 6864221.000000 ) AND ( ABMCU BETWEEN ' 1' AND ' 20' )
```

The user MG5700778 will not be able to insert or delete any records in the Address Book table. Users of inclusive row security need to be aware that the presence of a single record or a set of security records in the F00950 Security Workbench Table with all N values for one or more operations for a table/data dictionary combination will disallow that user from performing that particular operation on the table.

Scenario: Exclusive Row Security

The following identifies the required records that are created in Row Security for an example with exclusive row security. In this example, user MG5700778 is able to view records in the Address Book table (F0101) with a business unit value between 1 and 20 and between 51 and 70. However, that user is only allowed to update records in the Address Book table with a business unit value between 1 and 20. For exclusive row security the following four records are added for row security:

| User | Table | Data item | From value | Thru value | Add | Change | Delete | View | Alias |
|-----------|-------|------------|------------|--------------|-----|--------|--------|------|-------|
| MG5700778 | *ALL | CostCenter | 1 | 20 | N | Y | N | Y | MCU |
| MG5700778 | *ALL | CostCenter | 21 | 50 | N | N | N | N | MCU |
| MG5700778 | *ALL | CostCenter | 51 | 70 | N | N | N | Y | MCU |
| MG5700778 | *ALL | CostCenter | 71 | ZZZZ ZZZZ | N | N | N | N | MCU |

Selects performed against the F0101 table would resemble the following:

```
SELECT * FROM TESTDTA.F0101 WHERE ( ABMCU NOT BETWEEN ' 21' AND ' 50' AND ABMCU NOT BETWEEN ' 71' AND ' ZZZZZZZZ' ) ORDER BY ABAN8 ASC
```

Updates on the F0101 table would resemble the following:

```
UPDATE TESTDTA.F0101 SET ABALKY='MG5700778',ABTAX='456456456',ABALPH='John Doe',ABDC='JOHNDOE',ABMCU=' 1',ABSIC=' ',ABLNGP=' ',ABAT1='E', ABCM=' ',ABTAXC=' WHERE ( ABAN8 = 9999999.000000 ) AND ( ABMCU NOT BETWEEN ' 21' AND ' 50' AND ABMCU NOT BETWEEN ' 51' AND ' 70' AND ABMCU NOT BETWEEN ' 71' AND ' ZZZZZZZZ' )
```

Business Unit Security

JDE World supports business unit security for fields ending in MCU.

Role Relationship Support

EnterpriseOne replaces group relationships with role relationships. Instead of assigning users to a specific group profile, users are assigned specific roles that are globally defined through EnterpriseOne. The Role Relationship application establishes role relationships for applications, columns, and rows.

Roles can have overlapping privileges. For example, users belonging to different roles may need to perform common operations. Some general operations can be performed by all users (as is the case with *PUBLIC authority). In this situation, it would be inefficient to repeatedly specify these general operations for each new role that is created. Role hierarchies can be established to help alleviate this repetition. A role hierarchy defines roles that have unique attributes and that may contain other roles; that is, one role may implicitly include the operations associated with another role. Roles can be nested down two levels, and the ordering rules are recognized at both levels.

All roles assigned to a user adhere to the sequences as defined in EnterpriseOne. For example, assume the following roles are defined as:

```
RoleA  3
RoleB  2
      RoleF  1 (this role is nested under RoleB)
RoleC  4
```

Note: The underlined numbers indicate the defined role order.

If User1 is assigned to RoleA, RoleB, and RoleF, RoleF takes precedence over RoleB (which, in turn, takes precedence over RoleA). Ordering rules apply to all Showcase products supporting EnterpriseOne security.

In the above example, the role relationship rules for User1 would be:

```
User1
RoleF
RoleB
RoleA
*PUBLIC
```

In cases of conflicting role security privileges, the highest level user-assigned role takes precedence.

The following role relationship guidelines apply:

- Every role must have the appropriate environments defined. For example, if a user is assigned to RoleA, RoleB, and RoleF, all three roles must contain the same environment definitions.
- All roles assigned to a user must be set to be included in the *ALL role group.

SETJDEOWA Command

The SETJDEOWA command configures the server on the IBM i for JDE OneWorld and EnterpriseOne integration. The command has a number of parameters, each of which is discussed in detail in the following sections.

Note: SETJDEOWA requires that the IBM i interactive job CCSID is not set to 65535 (*HEX). SETJDEOWA includes the following commands:

OneWorld Release (JDEOWRLS)

Specifies the OneWorld or EnterpriseOne release installed on the IBM i.

JDE OneWorld System Library (JDEOWSYSLB)

Specifies the name of the IBM i library that contains the OneWorld security tables. When entering this value, pressing F4 provides a list of possible library names. You can specify to search the system for the JDE system library or explicitly specify the library name.

Use JDE OneWorld App Security (APPSECOPT)

Allows JDE OneWorld application security to function as table security. The application security parameter provides the options *DISABLE or *ENABLE (*DISABLE is the default setting). Enabling application security forces the honoring of any OneWorld application security for tables that the OneWorld administrator has set up for users and groups.

Run Authority Override (RUNAUTHOVR)

Controls how application security is applied for tables shared by multiple OneWorld applications. The Run Authority Override parameter only applies to the enabling of application security. If two OneWorld applications use the same database table, and the application security conflicts, where the user is authorized to one OneWorld application but not to the other, this option allows the administrator to select the default authorization for the application in question. If this parameter is set to *YES, the table is available to the user if the user has run authority for at least one application that uses the database table. If this parameter is set to *NO, the user is restricted from the database table if the user is restricted from at least one application that uses the database table. The default value is *YES.

Single JDE OneWorld Profile (SNGLUSRPRF)

The single OneWorld profile option allows you to enable or disable JDE OneWorld non-IBM i profile sign-on. When this parameter is set to *ENABLE, the user does not need an IBM i user profile with the same name. When this parameter is set to *DISABLE, the user must have an

IBM i user profile with the same name. The database user for the Showcase session is obtained from the OneWorld system security tables.

Authentication System Name (AUTHSYSNAM)

Defines the system used to validate OneWorld/EnterpriseOne application users. The special value of *LOCAL should be used if the local IBM i has the OneWorld/EnterpriseOne services. If specifying a remote IBM i, the remote system name must be the RDB name. The RDB name must also be accessible as a host name via TCP/IP.

Authentication User (AUTHSYSUSR)

Specifies the IBM i user profile used when connecting to the authentication system. The value must be *NONE for an authentication system of *LOCAL. The value must be a valid *ALLOBJ profile on the remote system when the authentication system is specified as an RDB name (other than *LOCAL).

Authentication Password (AUTHSYSPWD)

Specifies the IBM i user profile password used when connecting to the authentication system. The value must be blank for an authentication system of *LOCAL. The value must be the correct password when the authentication system is specified as an RDB name (other than *LOCAL).

JDE OneWorld Program Library (JDEOWPGMLB)

Specifies the name of the IBM i library containing OneWorld programs and service programs. When entering this value, pressing F4 provides possible library names.

Note: IBM Collaboration and Deployment Services Deployment Portal requires a true IBM i user profile and does not utilize the Showcase OneWorld Application User feature.

JDE OneWorld Remote System Lib (JDEOWRSYLB)

Specifies the name of the OneWorld/EnterpriseOne system library on the remote system. The OneWorld Remote System Library parameter is used only when the Authentication System is not *LOCAL. The special parameter value of *LCLSYSLB can be used to indicate that the remote system library name is the same as the local system library name.

JDE Multiple Library List Environments

Showcase applications retrieve library list environments that exist in the JDE environment and support multiple library lists for each OneWorld user profile. The following conditions apply:

- Users must have at least one defined library list environment for EnterpriseOne 9.0.
- Libraries cannot be defined in the environment master file for EnterpriseOne 9.0.

***Note:** Library lists are built dynamically (user interaction is not required).*

The following sections discuss library list options for the various JDE applications.

Selecting Group Library Lists for World

JDE must follow a specific order when selecting group library lists. This order differs for OneWorld and World users.

Currently, JDE World depends on the initial programs-J98INIT or J98INITA. The library list selection order 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 library list selection order for World users with J98INITA 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)

***Note:** Showcase Query, Showcase Report Writer, and IBM Collaboration and Deployment Services Deployment Portal will not function properly if the library list cannot be set, due to incorrect data, missing libraries, or the lack of a library list.*

Selecting Group Library Lists for OneWorld

JDE must follow a specific order when selecting group library lists. This order differs for OneWorld and World users.

The library list selection order 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)

Editing showcase.cfg to Use an Alternative F0094 File

Usually, the *F0094* file in JDE OneWorld contains the Master Library List data. JDE OneWorld integration with Showcase applications relies on the data in this file to set the library list upon connecting to the database through the web or 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 Showcase applications cannot determine the data source used for the library lists, a method has been provided that allows users to manually specify an alternative file (which must contain *F0094* data). This method involves duplicating the *F0094* file and editing the configuration file (*showcase.cfg*).

To set up the alternative configuration:

1. Find the JDE system library name by using the WRKOBJ command:

```
WRKOBJ OBJ(*ALL/F98OWSEC) OBJTYPE(*FILE)
```

The object and library name are displayed.

2. Because a duplicate of the *F0094* file is required, use the CRTDUPOBJ command, with the *F0094SPSS* filename in MYLIB and the JDE system library name *SYS7333* as examples:

```
CRTDUPOBJ OBJ(F0094) FROMLIB(SYS7333) OBJTYPE(*FILE)
TOLIB(MYLIB) NEWOBJ(F0094SPSS) DATA(*YES)
```

Note: The DATA (*YES) option duplicates the data. If you do not want to populate an empty file with the library list information, you can set the option to DATA (*NO).

3. The main configuration file (*showcase.cfg*) in the server library must contain two entries specifying the location of the alternative filename and library name. Use the IBM i command EDTF to edit the file.

Client access can be used for the same purpose. Following is the EDTF example, where SCSERVER is the server library name:

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

Add the following two lines to the configuration file:

```
showcasecorp.JDEOneWorldF0094File=F0094SPSS
showcasecorp.JDEOneWorldF0094Library=MYLIB
```

4. To activate this change, stop and restart the server. If IBM Collaboration and Deployment Services Deployment Portal is used, with Report Services enabled, you must also restart your application server.

Selecting Library Lists for EnterpriseOne via ADDE1LIBL

The Showcase 9 server builds the library list for a chosen environment based on the data source information retrieved from the EnterpriseOne installation.

Note: The library list is all-inclusive and includes environments derived from users, roles, and *PUBLIC.

The ADDE1LIBL command allows Showcase 9 to specify additional data libraries for a given EnterpriseOne environment. When an EnterpriseOne end user connects to the Showcase 9 server using Showcase Query client and selects a library list environment, any additional library list entries are added to the connection's library list (directly after the EnterpriseOne data library entry).

Command Syntax for Adding Library List Entries to an Environment

```
ADDE1LIBLE E1ENVNAME (EnterpriseOne_Environment_Name)
E1ENVLIBLE (MYDATALIB1 MYDATALIB2)
```

EnterpriseOne_Environment_Name must be an EnterpriseOne registered environment. Up to 240 library list names can be specified. The libraries need not exist at the time of running the command, but they must exist when end users connect and select the EnterpriseOne environment name.

If the command is run multiple times for the same environment name, all previously defined library list entries are replaced with the entries specified in the command.

Example: To add the library list entries MYDATALIB1 and MYDATALIB2 to the production environment PD900, use the following command:

```
ADDE1LIBLE E1ENVNAME (PD900) E1ENVLIBLE (MYDATALIB1 MYDATALIB2)
```

Command Syntax for Removing Library List Entries from an Environment

```
ADDE1LIBLE E1ENVNAME (EnterpriseOne_Environment_Name)
E1ENVLIBLE (*REMOVE)
```

All library list entries for the specified environment are removed.

Example: To remove all library list entries from the development environment DV900, use the following command:

```
ADDE1LIBLE E1ENVNAME (DV900) E1ENVLIBLE (*REMOVE)
```

Example before-and-after library list entries

The following table provides an Enterprise One example before-and-after production environment (PD900) library list when MYDATALIB1 and MYDATALIB2 are defined via the ADDE1LIBLE command:

| Before ADDE1LIBLE | After ADDE1LIBLE |
|----------------------|---------------------|
| QTEMP | QTEMP |
| PRODDTA | PRODDTA |
| PRODCTL | MYDATALIB1 |
| COPD900 | MYDATALIB2 |
| SY900 | PRODCTL |
| DD900 | COPD900 |
| OL900 | SY900 |
| QGPL | DD900 |
| | OL900 |
| | QGPL |

Default installation library names

| Library description | Default installation libraries (examples) EnterpriseOne 9.0 |
|----------------------------|--|
| Data Dictionary Library | DD900 |
| Object Librarian Library | OL900 |
| System Library | SY900 |
| Production Control Library | PRODCTL |
| Production Data Library | PRODDTA |

JDE User Authentication in IBM Collaboration and Deployment Services

The Showcase JDE OneWorld Application User Security Provider for IBM Collaboration and Deployment Services allows end users of IBM Collaboration and Deployment Services Deployment Manager and IBM Collaboration and Deployment Services Deployment Portal to login with an Oracle JDE Application User. This user is defined only in the JDE OneWorld/EnterpriseOne environment.

Configuration prerequisites

The prerequisites for the configuration are as follows:

1. Ensure that the IBM Collaboration and Deployment Services - Server Adapter for Showcase is installed into IBM Collaboration and Deployment Services.
2. Ensure that your Showcase Warehouse Manager Server is configured for your version of Oracle JDE OneWorld or EnterpriseOne and that Showcase Query and Showcase Report Writer users can connect with JDE Application Users.


Configuring the JDE Application User Security Provider

To configure the JDE Application User Security Provider, follow these steps:

1. Login to IBM Collaboration and Deployment Services Deployment Manager (either the installed client Server Administration or the web interface for configuration). See your IBM Collaboration and Deployment Services documentation for details.
2. Navigate to the Security Providers component.
3. Edit the settings for the JDE Application User Security Provider.

Figure 4-1

JDE Security Provider Configuration



4. In the IBM i Server box, enter the TCP/IP name (or IP address) of your system where the Showcase Warehouse Manager Server is installed.

5. In the Showcase Server Port box, enter the TCP/IP port number for your Warehouse Manager Server.
6. Select the Enable checkbox to enable the JDE Application Security Provider.
7. Click the X icon to close the dialog box and save your changes.

Configuring JDE EnterpriseOne 9.2

To configure JDE EnterpriseOne 9.2 for Showcase 9, follow these steps:

1. Issue the *Set JDE OneWorld/E1 Attributes* command - **SETJDEOWA** - and press **F4** to prompt.
2. Fill in the parameters for the following values:
 - JDE Release = ***ERP9.2** (use F4 for a list of valid values)
 - JDE System Library = **SY920**
 - Use JDE Application Security = ***ENABLE**
 - JDE Program Library = **E920SYS** (use F4 for a list of valid values)
3. Press Enter.

This creates 4 files in the IFS:

- /SCSERVER/proddata/config/jdbj.ini
- /SCSERVER/proddata/config/jdelog.properties
- /SCSERVER/proddata/config/ssinterop.ini
- /SCSERVER/proddata/config/JDEOW.cfg

Verify that the ssinterop.ini file includes correct information for JD Edwards, such as the port, Enterprise Server name, Security Server name and data source name:

4. Define system users and corresponding passwords (required for EnterpriseOne 9.2.0.5 and higher) by issuing the command: **SETJDESEC**.
5. On the Edit JDE Mapped Profiles screen, select the option to add (*ADD) and enter the User Profile **JDE**.

This creates the following file in the IFS that stores System User profiles with encrypted passwords.

- /SCSERVER/proddata/config/JDEOWSEC.cfg

It can be deleted and recreated if necessary, however, it is not possible to edit the file, since there's no way to decrypt or encrypt the passwords.

Use the SETJDESEC command with options *CHG or *DLT to edit the entries stored in this file.

EnterpriseOne 9.2.1.0 and higher additional step required

6. Edit the showcase.cfg file using the command shown, and set the Java version used by the Showcase server. This is not required for Tools Releases prior to 9.2.1.0.
EDTF '/SCSERVER/proddata/config/showcase.cfg'
7. In the editor, insert 2 blank lines near the top of the file (just under the 'Showcase Server Properties' section).

8. On one of the new blank lines add the following link (ensure there are no leading spaces):

```
showcasecorp.scsrvr.jvm.path=/QOpenSys/QIBM/ProdData/  
JavaVM/jdk70/64bit
```

Other Java paths that could be used **if supported by JD Edwards EnterpriseOne:**

```
/QOpenSys/QIBM/ProdData/JavaVM/jdk71/64bit
```

```
/QOpenSys/QIBM/ProdData/JavaVM/jdk80/64bit
```

Notes:

- Refer to Certifications on My Oracle Support for additional information on which versions of Java are supported for your Tools Release
 - 32-bit Java is not recommended for ShowCase applications for performance reasons.
9. Save the change by hitting **F3** twice. Restart the server in order for the change to take place:

SCSERVER/ENDSCSVR

SCSERVER/STRSCSVR

Server Maintenance

Using TCP/IP

The product software is used in a transmission control protocol/Internet protocol (TCP/IP) network environment. Client and server products communicate with each other using connections through TCP ports, and firewalls can be enabled to restrict access and improve security.

To change TCP port numbers after the installation, use the `CHGSCSVRA` command. To change firewall settings, use the `ADDSTCP` and `RMVSTCP` commands. These commands make it easier to migrate from a test environment to a production environment. See “Migrating Users to a New Server Library” on page 49 for more information.

Configuring TCP/IP with STRSCSVR

The Showcase Warehouse Manager Server TCP/IP settings can be configured during installation or after installation. Before configuring TCP/IP, verify that IBM i TCP/IP is configured and active. Run the `CHGSCSVRA` command to enable Warehouse Manager Server TCP/IP support.

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

- `CHGSCSVRA` changes subsystem information and TCP/IP-specific objects
- `STRSCSVR` starts the manager job
- `ENDSCSVR` ends 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. Each Showcase Warehouse Manager Server installation library is assigned a specific port. To enable the server for TCP/IP, you must assign a unique port number for each library. Run the IBM i command `WRKSRVTBLE` to view a list of the ports in use.

The command `ADDSTCP` assigns the port number 43419 by default. If this number already displays when you run the `WRKSRVTBLE` command, choose another port number. The next available number is specified by default.

Migrating Users to a New Server Library

It may become necessary to have your users change libraries. For example, you could install a new server version and must first verify it works in a test environment. In this case, you want users to stop using the old version and begin using the new one.

In a TCP/IP network, a user's client configuration employs the TCP/IP port to determine which Showcase Warehouse Manager Server 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. In combination with unique firewall configurations, users are automatically connected to the new library without any changes to their client configuration.

Configuring Firewall Settings with ADDSCTCP

Showcase Warehouse Manager Server firewalls can be configured at installation or after installation. After installation, use the ADDSCTCP and RMVSCTCP commands to configure the subsystem, TCP/IP ports, and the RMI host name used for the Warehouse Manager Server. Configurable ports include the main listening port, the firewall listening port, and the compressed data port. The ports are added to the IBM i service table and can be viewed with the WRKSRVTBLE command by looking for entries beginning with SCSEVER, SCFIREWALL, and SCRMIComp and ending with the name of the IBM Warehouse Manager Server library. The RMI host name is stored in an Showcase internal object.

Note: Use RMVSCTCP to remove existing configuration settings before assigning new values with ADDSCTCP.

Backward Compatibility with showcase.cfg

Firewall settings can also be entered manually in the configuration file `showcase.cfg`, but this method is not recommended. Showcase will check for firewall entries in `showcase.cfg` and use them if the firewall ports and RMI hostname are not defined elsewhere (at installation or with ADDSCTCP). Firewall entries in `showcase.cfg` should follow these formats:

```
showcasecorp.firewallport=<port_no>
showcasecorp.compression.firewallport=<port_no>
showcasecorp.rmi.server.hostname=<hostname>
```

Firewall settings entered manually in `showcase.cfg` also need to be removed manually if the ADDSCTCP command is used to configure the settings later. If the firewall configuration entries in `showcase.cfg` conflict with firewall values set with ADDSCTCP, the ADDSCTCP settings take precedent.

Note: The configuration file is located by default at `/<SCSEVER>/proddata/config/showcase.cfg`, where SCSEVER is the name of the Warehouse Manager Server installation library.

Checking Server Library Usage

Listing Connections to Relational Applications Only

The following procedures list active TCP/IP jobs to relational products, such as Showcase Query, Showcase Report Writer, Showcase Warehouse Manager Client, and Showcase Warehouse Builder. To list all active application connections, see Listing All Connections to All Applications on page 51.

Run the following command for Showcase Warehouse Manager Server libraries:

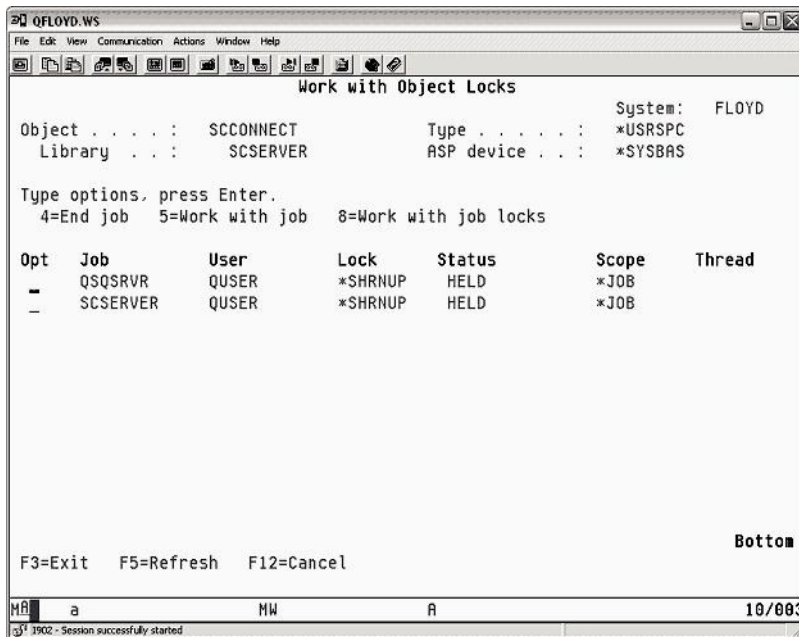
```
WRKOBJLCK OBJ(SCSEVER/SCCONNECT) OBJTYPE(*USRSPC)
```

where SCSEVER is the name of the library you want to check.

If TCP/IP is active, a screen similar to Figure 5-1 shows the TCP/IP job for the library (job name SCSEVER, which is the same as the library name) and may show TCP/IP connections.

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 that no user is connected. From the user profile name, you should be able to determine who is using the product. Alternatively, you can use the Warehouse Manager Client software to view active connections.

Listing All Connections to All Applications

This method lists active connections to all applications.

1. Open the Showcase Warehouse Manager Client application and connect to a data source.
2. Double-click the data source for which you want to view connections.
3. From the Warehouse Manager menu, choose Applications.
4. From the Manage menu, choose Active Connections.

Ending TCP/IP (Stopping the Server)

Before ending TCP/IP, verify no one is using the library (refer to "Checking Server Library Usage" on page 50). If you see the library name in the job lock list, TCP/IP is enabled.

From the library in which you want to end TCP/IP, run the command `ENDSCSVR`. This command ends TCP/IP clients and stops the TCP/IP server job for the specified library.

Uninstalling a Server Library

You may occasionally need to uninstall server libraries on your IBM i system (such as test libraries).

To uninstall server libraries:

1. Log on as QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).
2. Back up any Showcase Warehouse Manager Server settings you want to save. See Backing Up Server Information on page 53 for instructions.
3. Verify that no users are connected to the library. See Checking Server Library Usage on page 50.
4. Restore the UNINSTSC utility from the UNINSTSC save file that exists in the server library by entering the following command:

```
RSTOBJ OBJ(*ALL) SAVLIB(QTEMP) DEV(*SAVF) SAVF(SCSERVER/  
UNINSTSC) MBROPT(*ALL) ALWOBJDIF(*ALL) RSTLIB(QTEMP)
```

where SCSERVER is the name of the Warehouse Manager Server library to be uninstalled.

Note: This restores three objects into QTEMP.

5. Enter the command QTEMP/UNINSTSC and press Enter. This command cannot be submitted to batch mode.
6. Enter the name of the Warehouse Manager Server library you want to uninstall and then press Enter to confirm.

Reloading the Server Configuration

You may occasionally need to reload the server configuration without stopping and restarting the server. The RELODSCCFG command allows the majority of the Showcase configuration and JDE OneWorld/EnterpriseOne configuration to be re-initialized without requiring a server restart. The command's primary function is to re-initialize the properties associated with the Java logging feature. Secondly, the JD Edwards information can be re-initialized if the SETJDEOWA command is run prior to the RELODSCCFG command.

The RELODSCCFG command does not affect any parameters directly associated with the Java Virtual Machine. Such properties include the standard input, standard output, and verbose options. Changes for these properties require a Showcase Warehouse Manager Server restart, since they are only able to be set prior to the start of the Java Virtual Machine used in the Warehouse Manager Server.

To reload the server configuration:

Note: Running the RELODSCCFG command has no effect when Warehouse Manager Server is not running.

1. Log on to the IBM i with a user profile that has *ALLOBJ authority.
2. Add the Warehouse Manager Server source library to your IBM i library list and ensure you have only one Warehouse Manager Server library in your library list. Run the command:

```
ADDLIBLE SCSERVER
```

where SCSERVER is the name of the Warehouse Manager Server library.

3. Enter the RELODSCCFG command and press Enter to reload the server configuration.
4. Sign off of the IBM i.

Backing Up Server Information

Due to the nature in which Showcase Warehouse Manager Server is installed, it is recommended that you back up your installation library with SAVSRVRINF rather than using conventional means. Following is the series of procedures used when backing up a server library.

Saving Server Information

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

- Showcase Warehouse Manager Client (server options such as IBM i library access and default public authority to IBM i libraries and tables; security and resource settings, aliases, and column attributes; and application security)
- Showcase Query (data views)
- Showcase Warehouse Builder (data distribution sets and definitions)
- IBM Collaboration and Deployment Services (as a repository to store Query files)

The Warehouse Manager Client, Query, Warehouse Builder, and Showcase Report Writer applications store information on the server on the IBM i. 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 the two aforementioned commands, is also useful for moving data from one server library to another. See Moving Data from a Test to a Production Environment on page 29 and Restoring Data on page 31 for recommendations on moving data from a test to a production environment and backing up data.

***Note:** It is not recommended to use the SAVSRVRINF and RSTSRVRINF commands for the purposes of backing up or restoring an IBM Collaboration and Deployment Services repository. This option is available for migration purposes. Please refer to the IBM Collaboration and Deployment Services documentation for normal backup and recovery procedures for the repository*

To back up information that applications store on the server:

1. Log on to the IBM i with a user profile that has *ALLOBJ authority.
2. Add the Showcase Warehouse Manager Server source library to your IBM i library list and ensure you have only one Warehouse Manager Server library in your library list. Run the command:

ADDLIBLE SCSERVER

where SCSERVER is the name of the Warehouse Manager Server library.

3. Before saving server information, stop the server with the ENDSCSVR command. For more information, see Chapter 2.
4. Enter the SAVSRVRINF command and press F4.
5. At the Library prompt, enter the server library name from which server information should be saved. This library is typically SCSERVER.
6. At the Save file prompt, enter the save filename and the IBM i library in which you want to create the save file. Press Enter.
7. At the Type of Information to Save prompt, enter one of the following options. Press F1 to display help.

- *ALL to save all information described in the following bullets. You will be prompted to enter the server library name where IBM Collaboration and Deployment Services Repository is installed on the IBM i and the IBM Collaboration and Deployment Services administrative user ID and password.
- *ADMIN to save only Warehouse Manager Server administration information (such as security settings and aliases).
- *PES to save only IBM Collaboration and Deployment Services information (such as queries, reports, scheduled jobs, and other files saved to the repository). You will be prompted to enter the server library name where IBM Collaboration and Deployment Services Repository is installed on the IBM i and the IBM Collaboration and Deployment Services administrative user ID and password.
- *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.

8. Press F10 to view additional prompts for the SAVSRVRINF command. These prompts need not be changed under most circumstances. The prompts are described below:
 - At the Target Release prompt, enter *PRV if you intend to restore saved information on the previous IBM i release. Leave the default of *CURRENT if you intend to restore saved information on the same or a later IBM i release.
 - 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. This allows you to change a physical file without deleting dependent views.
9. Press Enter to save the server information.
10. Sign off of the IBM i.

Restoring Server Information

If you plan to restore server information to a new server library, first follow the procedures in the Showcase 9 Installation Guide to install the new server.

Restoring Server Information

Use the RSTSRVRINF command to restore server information saved with the SAVSRVRINF command. RSTSRVRINF allows you to restore Showcase Warehouse Manager Server information, Showcase Warehouse Builder information, and some IBM Collaboration and Deployment Services information.

If you choose to restore Warehouse Manager Server and/or IBM Collaboration and Deployment Services information, all information is restored. If you choose to restore Warehouse Builder information, you have the option of restoring all information or restoring specific Warehouse Builder sets or definitions. This can be useful when, for example, you want to restore a corrupt or inaccurate Warehouse Builder set to a previous state. Instead of restoring all Warehouse Builder sets, you can choose to simply restore the set in question.

To restore server information saved with the SAVSRVRINF command:

1. Log on to the IBM i with a user profile that has *ALLOBJ authority.
2. Distribute your save file to the target system (if it's different from the first system).
3. Add the Warehouse Manager Server source library to your IBM i library list and ensure that you have only one Warehouse Manager Server library in your library list. Run the command:

ADDLIBLE SCSERVER

where SCSERVER is the name of the Warehouse Manager Server library.

4. Before restoring server information, stop the server with the ENDSCSVR command. For more information, see Chapter 2. *Note:* When restoring IBM Collaboration and Deployment Services information, the IBM Collaboration and Deployment Services server must be running.
5. Enter the RSTSRVRINF command and press F4.
6. At the Save file prompt, enter the save filename and the IBM i library in which you created the save file. Press Enter.
7. Choose whether or not to restore 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 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 Showcase Warehouse Manager Client application. On the Manage menu, choose Clean Up Showcase Catalogs, and then click View Results to display the results.*

8. If you specified *YES at the Restore Warehouse Manager Info prompt, the Duplicate Record Option prompt displays. Choose whether to keep or replace duplicate records by entering the following:
 - *KEEP to keep duplicate records.
 - *REPLACE to replace duplicate records with those found in the restore object.
9. Choose whether or not to restore Warehouse Builder information. Press Enter. If *YES, the Duplicate Record Option, Restore Logs, Restore Scheduled Sets, and WB objects to be restored prompts display. When the Restore Warehouse Builder Info prompt is set to *PRINT, values for the Restore Warehouse Manager Info, Restore Repository, and Restore Data View Information prompts must be *NO.
10. If you specified *YES at the Restore Warehouse Builder Info prompt, the Duplicate Record Option prompt displays. Choose to keep or replace duplicate records by entering the following:
 - *KEEP to keep duplicate records.
 - *REPLACE to replace duplicate records with those found in the restore object.
11. If you specified *YES at the Restore Warehouse Builder Info prompt, the Restore Logs prompt displays. Choose whether or not to restore the logging and message information from previous runs of the data distribution sets and definitions that you saved.
 - *NO does not restore logging and message information from previous runs.

- *YES restores logging and message information from previous runs.
12. If you specified *YES at the Restore Warehouse Builder Info prompt, the Restore Scheduled Sets prompt displays.
- *NO does not restore Warehouse Builder schedule information. The information includes the date, time, and frequency for each scheduled instance.
 - *YES restores Warehouse Builder schedule information. The information includes the date, time, and frequency for each scheduled instance.
13. If you specified *YES at the Restore Warehouse Builder Info prompt, the WB Objects to be Restored prompt displays. Choose one of the following options:
- *ALL restores all Warehouse Builder information. This is the default setting.
 - *SET restores specific Warehouse Builder sets. All objects contained in the specified sets are restored.
 - *DEF restores specific Warehouse Builder definitions. When a selected definition is restored, it becomes uncontained unless it replaces an existing contained definition. When a definition is contained in an existing set and a restored set (with *REPLACE option) has the same definition name, the new set is created and the existing definition is overwritten. This overwritten definition now exists in two sets. If a definition is contained in an existing set and a restored set (with *KEEP option) has the same definition name, the new set is created and the existing definition is kept, but not attached to the new set.

***Note:** Restoring individual Warehouse Builder sets or definitions is applicable only when the Restore Repository prompt is set to *NO.*

14. If you specified *SET or *DEF at the WB Objects to be Restored prompt, the Object Names to Restore prompt displays, allowing you to specify which Warehouse Builder sets or definitions are restored. Use the following guidelines when entering set or definition names:
- Object names are case sensitive.
 - Typing (+) allows you to view more input value lines.
 - Object names do not need to be enclosed in single quotes unless they contain special characters (\$, @, %, and so on).
 - A single quote in the object name is represented by two single quotes, while the whole string must be enclosed in single quotes. For example, you would enter 'O"Sullivan' for O'Sullivan.
 - Use an asterisk (*) when you do not know the characters in the object's name or you want to restore objects with similar names. An asterisk (*) stands for an unknown string of zero or more characters. For example, ABC* restores a list of all objects that begin with the characters ABC (ABC, ABCD, or ABCTEST).
 - The maximum number of listed objects is 50.

Keep the following points in mind when restoring individual Warehouse Builder sets or definitions:

- To determine which specific set and/or definition names are contained in a save file, set the Restore Warehouse Builder Info prompt to *PRINT. This produces a structured print file that provides the save file contents. When the Restore Warehouse Builder Info prompt is set

to *PRINT, values for the Restore Warehouse Manager Info, Restore Repository, and Restore Data View Information prompts must be *NO.

- If a distribution is sent to a different server, the user is responsible for creating the target tables on that system. The recommended method for creating target tables is to open the definition using Warehouse Builder client and resave the table.
- Restoration of individual Warehouse Builder sets and/or definitions is currently not supported on double-byte character set (DBCS) IBM i servers.

15. Choose whether or not to restore IBM Collaboration and Deployment Services information (such as queries, reports, scheduled jobs, and other files saved to the IBM Collaboration and Deployment Services Repository). Press Enter. If *YES, you will be prompted to enter the server library name where the repository is installed on the IBM i and the IBM Collaboration and Deployment Services administrative user ID and password. The user ID credentials must be for a local user repository administrator, not an IBM i administrator.

Important Consideration: *If you are upgrading from one version of IBM Collaboration and Deployment Services to another, it is recommended that you follow the documented procedure from IBM, not this command. Please contact Help/Systems Technical Support before using RSTS-RVRINF to restore a repository.*

16. Choose whether or not to restore Data View information.

17. Press F10 to see additional prompts for the RSTSRVRINF command. The following prompts need not be changed under most circumstances:

In the Catalog File to be Processed prompt, the default *ALL restores all catalog files that contain application administration information. These files contain application resource settings, Warehouse Manager Server security settings, aliases, and column attributes. Alternatively, you can specify a single catalog file to restore. Press F1 at this prompt to display the help, which lists the catalog files and their contents.

18. Press Enter to restore the server information. Several status messages display.

19. A screen with additional configuration steps displays. Read this information carefully, especially if you restored Warehouse Builder, Warehouse Manager, Showcase 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, it is recommended that you use Warehouse Builder to make any necessary changes, to create target tables, and to reschedule jobs.

MAPDIST Utility

The MAPDIST utility updates Showcase Warehouse Builder definitions by mapping the source or target data source to a new data source. The utility was created to resolve the following RSTS-RVRINF limitations:

- Enable Warehouse Builder to properly set or change the source and target connections.
- Provide the option of saving a definition with *LOCAL as the RDB name. This forces the local RDB to be used even if it is restored on a different system.
- Eliminates the error received in Warehouse Builder definitions when referenced SCSERVER libraries are deleted (and the definition is utilizing JD Edwards code). The error occurs because the DEFCTRL table retains the name of the Showcase Warehouse Manager Server library the definition was created with, not the library the definition is currently initiating from (the referenced library no longer exists on the IBM i).

User Interface

The MAPDIST command is included in all Warehouse Manager Server 9.0 libraries.

Figure 5-2

Map Warehouse Builder definitions (MAPDIST) utility

Map WB definitions (MAPDIST)

Type choices, press Enter.

| | | |
|---------------------------------|-------|-------------------------|
| Server type to map | _____ | *SOURCE, *TARGET, *BOTH |
| Old RDB name | _____ | |
| Old SC library name | _____ | Name |
| New RDB name | _____ | |
| New SC library name | _____ | Name |
| New ODBC data source name . . . | _____ | |

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display

F24=More keys

Bottom

Parameters

MAPDIST consists of the following parameters:

Server type to map. The server type to be mapped. Possible values include:

- *SOURCE. Maps the definition's source data source.
- *TARGET. Maps the definition's target data source.
- *BOTH. Maps the definition's source and target data source.

Old RDB name. The original name of the RDB to be mapped.

Old SC library name. The original name of the server library to be mapped.

New RDB name. The new RDB name to be mapped. Using *SAME retains the current RDB name. The user also can specify *LOCAL, resulting in local RDB usage (even after it is restored in a different system).

New SC library name. The name of the new server library to be mapped.

New ODBC data source name. The user-defined Showcase data source name. This is the data source name the Warehouse Builder client application uses when connecting to the IBM i.

Notes:

- All parameters are required.

- The commands cannot run while the server is running (the user receives an error indicating that the command cannot run).
- A warning message is logged when the old RDB or the old server library is not used by any existing Warehouse Builder definition.
- If the source RDB is changed, source tables must exist in the new system for the definition to work.
- If the target RDB is changed, the user is responsible for creating the target tables in that system. To create the target table, the user must open the definition with Warehouse Builder client and resave.

Detecting and Removing Obsolete Information

Use Showcase Warehouse Manager Client's Alias Manager and Security Manager to 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 Showcase Query to create data views, Warehouse Manager Client 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 Client information would then reference objects that no longer exist.

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

Restricting Access to a Server Library

It may be necessary to place a 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, allowing users to reconnect.

Applying the Restricted State

To place a server library in a restricted state:

1. Log on as QSECOFR or equivalent authority (*ALLOBJ, *IOSYSCFG, *SAVSYS, *JOBCTL, and *SECADM special authorities are needed).
2. Add the Showcase Warehouse Manager Server library to your IBM i library list and ensure that only one Warehouse Manager Server library is in your library list. Run the command:

ADDLIBLE SCSERVER

where SCSERVER is the name of the Warehouse Manager Server library.

3. Change the server status by entering the command CHGSRVRSTS and pressing F4.
4. At the Library prompt, enter the name of the Warehouse Manager Server library that you want to put in a restricted state. This library is typically SCSERVER.

5. 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.*

6. Press Enter.
7. At the Force Server Restricted State prompt, choose whether or not to force the server library into a restricted state:
 - *NO to attempt to place the server library in a restricted state. However, if users are connected to the server library, the attempt fails.
 - *YES to force all connections to the server library to end (except your current job and TCP/IP master job). This ends the jobs of any users who are currently connected to the server library.
8. 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 is 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 is not placed in a restricted state. In this case, you need to retry until the restricted state is achieved.

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

Removing the Restricted State

To release a server library from a restricted state:

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

ADDLIBLE SCSERVER

where SCSERVER is the name of the Warehouse Manager Server library.

3. Change the server status by entering the command CHGSRVRSTS and pressing F4.
4. At the Library prompt, enter the server library name you want to release from a restricted state.
5. At the Restrict Server Library prompt, enter *NO to allow users access to the server library.

6. Press Enter to remove the restricted state.

Unapplying Server Patches

With Showcase 9, administrators can run the `RMVSCPTF` command to roll back (unapply) the most recent Showcase Warehouse Manager Server patches. This command contains no parameters, but the server should be stopped prior to execution. Patches are rolled back to the prior patch level; the command restores the versions of all objects that existed before the latest patch was applied and removes all objects that did not exist prior to the latest patch. The command can roll back only one patch at a time but can be run in succession to roll back to the installation release level.

Note: The patch rollbacks rely on updated patch information in `SCPTFHST` (SQL table) and the availability of the backup save files and user space (created when patches are applied) in the Warehouse Manager Server installation library. If any of these is not found when the patch rollback is run, the rollback will not continue. To identify the difference between an original unaltered Warehouse Manager Server 9.0 installation from one in which patches were installed and subsequently removed, check the `SCPTFLOG` journal. If there are no records of patches being applied, then it is the original installation.

Working with Server Logs

Showcase 9 switches to standardized Apache LOG4J logging, which supports rollover logging among other features. Previously, the server logs could expand without limit, hampering the ability of users to manage the files and troubleshoot problems. In Showcase 9, logging properties are set in the `showcase.cfg` configuration file, located by default in the `/<SCSERVER>/ProdData/config/` folder, where `SCSERVER` is the name of the Showcase Warehouse Manager Server installation library.

Administrators can enable a rollover size option (`RolloverType = SIZE`) and specify a maximum log file size (`MaxFileSize`) so when the limitation is reached, a new log file is created and logging continues in the new file. Administrators can also enable a rollover date option (`RolloverType = DATE`), and each log file will roll over to a new file once the system date changes to the next day. With `MaxBackupIndex`, administrators can specify how many backup log files are kept before the oldest is erased. If this value is set to zero, then there will be no backup files and the log file will be truncated when it reaches `MaxFileSize`.

To enable logging, open the `showcase.cfg` configuration file with a text editor and set the following parameters:

Table 5-1
Log settings in `showcase.cfg`

| Property | Example |
|---|--|
| <code>traceService</code> | <code>showcasecorp.traceService=ALL</code> |
| <code>traceLevel</code> | <code>showcasecorp.traceLevel=ALL</code> |
| <code>traceUser</code> | <code>showcasecorp.traceUser=ALL</code> |
| <code>traceLog</code> | <code>showcasecorp.traceLog=scserver.log</code> |
| <code>traceRolloverType</code> | <code>showcasecorp.traceRolloverType=DATE</code> <code>showcasecorp.traceRolloverType=SIZE</code> |
| <code>traceMaxFileSize</code> (minimum of 50MB) | <code>showcasecorp.traceMaxFileSize=50MB</code> |
| <code>traceMaxBackupIndex</code> | <code>showcasecorp.traceMaxBackupIndex=10</code> |

Insert a hash mark (#) at the beginning of a line to comment out any log settings. Remove the hash mark to enable the log setting. In general, performance is decreased when logging of any type is enabled. However, the new rollover logging feature will not cause performance to degrade significantly.

***Note:** The rollover logging settings do not apply to the `scsserver.os400.stderr` and `scsserver.os400.stdout` logs, nor does it apply if the property `showcasecorp.traceLog` has a value of `JOBLOG`, as these entries are written to the IBM i job log.*

Working with Journals

Showcase Warehouse Manager Server creates four journals upon installation. These journals are used to maintain data integrity when catalog files are updated. The journals are named `QSQJRN`, `RSCSPCTLJ`, and `RSCSCCTLJ`, and `SCPTFLOG`.

The main concern with journal files is the space consumed by the journal receivers. Use one of the following approaches to ensure the journal receivers do not consume too much space:

- Monitor the size of the receivers and issue the IBM i command `CHGJRN` when they become 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. See your IBM i documentation for details on the `CHGJRN` command.
- Let the system maintain the journal receivers. See your IBM i documentation for changing the journal receivers and the journals for a user-defined threshold.

***Note:** Removing journal receivers limits the ability to determine the user(s) who have made changes to Showcase security, alias, and Showcase Warehouse Builder sets and definitions because the journal receivers contain the history of the changes to the Showcase control tables.*

Single Sign-On Configuration (SSO)

Single sign-on (SSO) is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. In order to use Showcase single sign-on, the following prerequisites must be met:

- The system of Showcase clients must belong to the company Intranet domain and a member of that domain must be logged into the client machine.
- IBM Enterprise Identity Mapping (EIM) must be configured and active on the IBM i where the Showcase Warehouse Manager Server is installed.
- The Warehouse Manager Server must be configured for the Intranet domain and the default EIM instance on the IBM i.
- The ShowCaseWarehouse Manager Server data source for the Warehouse Manager Server must be SSO-enabled.

Enterprise Identity Mapping (EIM) Configuration

Enterprise Identity Mapping (EIM) for the IBM i platform is the i5/OS® implementation of an IBM infrastructure that allows administrators and application developers to solve the problem of managing multiple user registries across their enterprise. Most network enterprises face the problem of multiple user registries, which require each person or entity within the enterprise to have a user identity in each registry. The need for multiple user registries quickly grows into a large administrative problem that affects users, administrators, and application developers. EIM enables inexpensive solutions

for easier management of multiple user registries and user identities in your enterprise. EIM allows you to create a system of identity mappings, called associations, between the various user identities in various user registries for a person in your enterprise. EIM also provides a common set of APIs that can be used across platforms to develop applications that can use the identity mappings that you create to look up the relationships between user identities. In addition, you can use EIM in conjunction with network authentication service, the i5/OS implementation of Kerberos, to provide a single sign-on environment. You can configure and manage EIM through iSeries Navigator, the IBM i graphical user interface. The IBM i platform uses EIM to enable i5/OS interfaces to authenticate users by means of network authentication service. Applications, as well as i5/OS, can accept Kerberos tickets and use EIM to find the user profile that represents the same person as the Kerberos ticket represents.

For more information about configuring Enterprise Identity Mapping for the IBM i platform, see the documentation at the IBM website.

Showcase Warehouse Manager Server Configuration

The Showcase Warehouse Manager Server needs the information on the Kerberos domain and the EIM credentials. The Kerberos Realm, KDC Server (Domain Controller), and the EIM administrative user and password are required when enabling Showcase SSO. The CFGSCSSO command in the Showcase library must be used to configure these settings.

Enabling SSO with the CFGSCSSO command

1. Add the Warehouse Manager Server library to your library list (with ADDLIB).
2. Type in the CFGSCSSO command and press F4 (instead of Enter).
3. Press F9 to show all required parameters.
4. Configure the settings, using the following information as a guide:

Enable Showcase SSO. This prompt allows you to enable or disable single sign-on for the Warehouse Manager Server. The possible values are *YES (single sign-on is enabled) and *NO (single sign-on is disabled).

Kerberos Realm. Enter the name of your Kerberos Realm. This is typically the name of your domain, in capital letters. Contact your network administrator to obtain the correct value for this parameter.

Kerberos KDC Server. Enter the name of your Kerberos Key Distribution Center (KDC). The KDC is typically installed on your Domain Controller. Contact your network administrator to obtain the correct value for this parameter.

EIM System. Enter the TCP/IP name of the IBM i where your Enterprise Identity Mapping (EIM) instance is configured and running. Contact your IBM i administrator to obtain the correct value for this parameter.

EIM Domain Name. Enter the name for your Enterprise Identity Mapping instance running on the IBM i machine. Contact your IBM i administrator to obtain the correct value for this parameter.

EIM LDAP Port. Enter the port for your Enterprise Identity Mapping LDAP Server (IBM Directory Server) running on the IBM i machine. Contact your IBM i administrator to obtain the correct value for this parameter.

EIM Administrative User. Enter the administrative user name for your local Enterprise Identity Mapping instance running on the IBM i machine. Contact your IBM i administrator to obtain the correct value for this parameter.

EIM Administrative Password. Enter the administrative password for your local Enterprise Identity Mapping instance running on the IBM i machine. Contact your IBM i administrator to obtain the correct value for this parameter.

EIM OS/400 Registry. Enter the name of the Enterprise Identity Mapping OS/400 registry. Contact your IBM i administrator to obtain the correct value for this parameter.

Once you have configured the single sign-on settings, restart the Warehouse Manager Server to pick up the changes.

Showcase IBM i Data Source Configuration

The Showcase IBM i Driver data source must be enabled for single sign-on. In the data source configuration dialog box, ensure that the checkbox for `Enable Single Sign-On` is selected.

Showcase Warehouse Manager Server and IBM i Work Management

This section describes how Showcase Warehouse Manager Server is influenced by IBM i work management controls and provides information for adjusting these controls. Before reading this section, you should be familiar with IBM i work management concepts. For information on IBM i work management, see the IBM i documentation.

Work Management Controls

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

In general, you should not need to adjust the work management behavior of Warehouse Manager Server. This information gives you a basic understanding of how Warehouse Manager Server 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 Showcase Warehouse Manager Server job class can be used to control the following attributes with minimum impact to other jobs:

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

Other attributes can also be controlled through the class description, such as time slice.

Job run priority. Used to control Warehouse Manager Server's run priority relative to other system jobs. If this is the only attribute you want to control, using the Resource Settings feature in Showcase Warehouse Manager Client is an easier alternative.

Maximum temporary storage. Used to control the amount of storage a job uses at runtime. Temporary storage is the storage required for Warehouse Manager Server and system programs at runtime, 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, the Warehouse Manager Server job allows an instruction to finish running. Typically, this is the time Warehouse Manager Server 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 allows Warehouse Manager Server to return record and object lock errors more quickly. Lengthening the wait time allows Warehouse Manager Server 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 because lock conditions may cause the PC to appear active 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 IBM i is a portion of the main memory allocated for a set of jobs. All jobs using a storage pool compete for the resource but do not compete against jobs using different storage pools. Storage pools are 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

Be careful when creating storage pools because the memory is subdivided and is not 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 that Showcase Warehouse Manager Server 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 page 16.

The subsystem the server job uses is determined as follows:

TCP/IP. TCP/IP uses the subsystem specified during the Warehouse Manager Server installation or with the CHGSCSVRA command. By default, this is either QCMN or QBASE on most systems. You must have multiple Warehouse Manager Server libraries installed if you need different PC workstations in different subsystems.

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

Routing Entries

Use routing entries to choose the job class used by Showcase Warehouse Manager Server 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 the product was installed or when the CHGSCSVRA command was issued. By default, the product uses the QWCPCSUP job class for TCP/IP connections.

To change the job class used by TCP/IP, use the CHGRTGE command.

The IBM i also manages TCP/IP. These jobs simply start the associated client jobs when a connection is made to the IBM i. The routing data are as follows:

Table 5-2
IBM i TCP 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 the job description is a specified batch job schedule, the routing data are 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 controls the initial library list used by the server, job log logging options, and job accounting.

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

Batch queries. For batch queries, the job description is 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

Showcase Warehouse Manager Server always runs using the user profile specified at logon when connecting to the IBM i.

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

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

Use one of the following techniques to identify a job by user. 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 is preserved for all users except *PUBLIC. *PUBLIC is unconditionally set to *EXCLUDE or *USE as described in the following tables.

Existing data administrators do not have authority to access the commands. You can:

- Grant authority to those users by employing the IBM i security management commands
- Use Showcase Warehouse Manager Client to grant the data administrator privilege

The commands in Table 5-3 are shipped with *PUBLIC authority set to *USE.

Table 5-3

*Commands shipped as *PUBLIC = *USE*

| Command | Description |
|-------------|--|
| DSPSCINF | Display Showcase Warehouse Manager Server tips |
| SBMDD | Submit Showcase Warehouse Builder distribution set |
| STRDD | Start Warehouse Builder distribution set |
| STRRPTEEXEC | Start job executor |

When granting a user data administrator privileges with Warehouse Manager Client, that user must be given IBM i *CHANGE authority to access the Warehouse Manager Server commands in Table 5-4. If the user is not given *CHANGE authority, some administration functions are not allowed.

Table 5-4

*Data administrator privilege = *CHANGE commands*

| Command | Description |
|------------|----------------------------------|
| CHGSRVRDFT | Change server default attributes |
| CHGSRVRSTS | Change server status |
| CHKSRVRCAT | Check server catalog information |
| SETLIBACC | Set library access |

User Exit Programs

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

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 the program before it runs a file transfer, virtual printer, or message function. Warehouse Manager uses the same network attribute and program parameters as Ware-

house Manager Server Access before it prepares each SQL statement or before submitting a batch query.

Following are descriptions of the two parameters passed to the user exit program when using Warehouse Manager or Warehouse Manager Server Access.

- 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 "I"). The request is rejected if this parameter is set to any other value.
- Parameter 2: This parameter contains information regarding the request. Table 5-5 describes the structure used by the product and is identical to the Warehouse Manager Server Access File Transfer structure.

Table 5-5

Structure of the second parameter passed to a user exit program (see footnote 1)

| 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 Warehouse Manager Server SQL. |
| Record format name | Character 10 | Contains the record format name. The product 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 Warehouse Manager Server Access 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 client application. Otherwise, the name is the client application name specified by the PC on the 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.

Unicode Considerations

The user exit program passes parameters in EBCDIC. Beginning with Showcase 7.1, if the SQL statement contains Unicode data that cannot be converted to EBCDIC, then either the exit program is called with substitution characters or the SQL statement is rejected.

The server is a multithreaded program. Exit programs that are not thread safe should be configured to run in a separate process. An exit program can be configured to run in a separate process using the MLTTHDACN (*NORUN) option when using ADDEXITPGM to register the exit program. This will

cause the server to run the exit program under a separate process using the connecting user's credentials. Optimum performance is achieved when running exit programs in the server process.

Distributed Relational Database Architecture (DRDA)

DRDA is the IBM implementation for database operations on remote DB2s. The correct operation of Showcase Warehouse Builder requires that connections to remote Showcase Warehouse Manager Server databases are properly set up. 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 Warehouse Manager Server database to another.

To ensure that remote databases are configured properly:

1. Enter the command `WRKRDBDIRE` and press Enter.

There should be a list of remote database names corresponding to other Warehouse Manager Server systems. If the required systems are present, the following step can be skipped.

2. Use the `ADDRDBDIRE` command to add any required remote Warehouse Manager Server DB2 databases. See your Warehouse Manager Server documentation for more information.

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.

Configure Showcase 9 Server to Reject Requests from Down-level Clients

This configuration option allows the Showcase 9 server to reject requests from down-level Client versions. This feature is available in server version 9.4.0.2 client version 9.1.0.13.

Follow the steps below to add an entry to the server configuration file, **showcase.cfg**, that will only accept connections from Showcase 9.1.0.13 or higher client versions.

Steps:

1. Edit the `showcase.cfg` file with the following command (where the Showcase install library is `SCSERVER`):

```
edtf '/scserver/proddata/config/showcase.cfg'
```

2. Insert a new line above the stanza - `"/ Properties must be added before the below stanza"` (empty lines above and below the new line are allowed)
3. Paste the following onto the new line:

```
minimumClientJdbcLevel=9.4.0.2
```

4. Restart the server for the change to take effect.

Down-level clients will receive a message, "Client software version is not compatible with the selected server".

The default value is to allow all clients. Therefore, to disable this feature, either remove the line entirely or change the value to `*ANY` like so:

```
minimumClientJdbcLevel=*ANY
```


Understanding Report Services

Report Services for IBM Collaboration and Deployment Services

IBM® Collaboration and Deployment Services uses Report Services to run reports and publish the results, in different formats, without requiring a PC such as an Enterprise Server. Report Services operates in a platform-independent environment, and provides a mechanism to load and interpret database queries and reports, retrieve data from various data sources, perform data manipulation such as calculations, and format the output into HTML, Microsoft Excel, PDF, or delimited text.

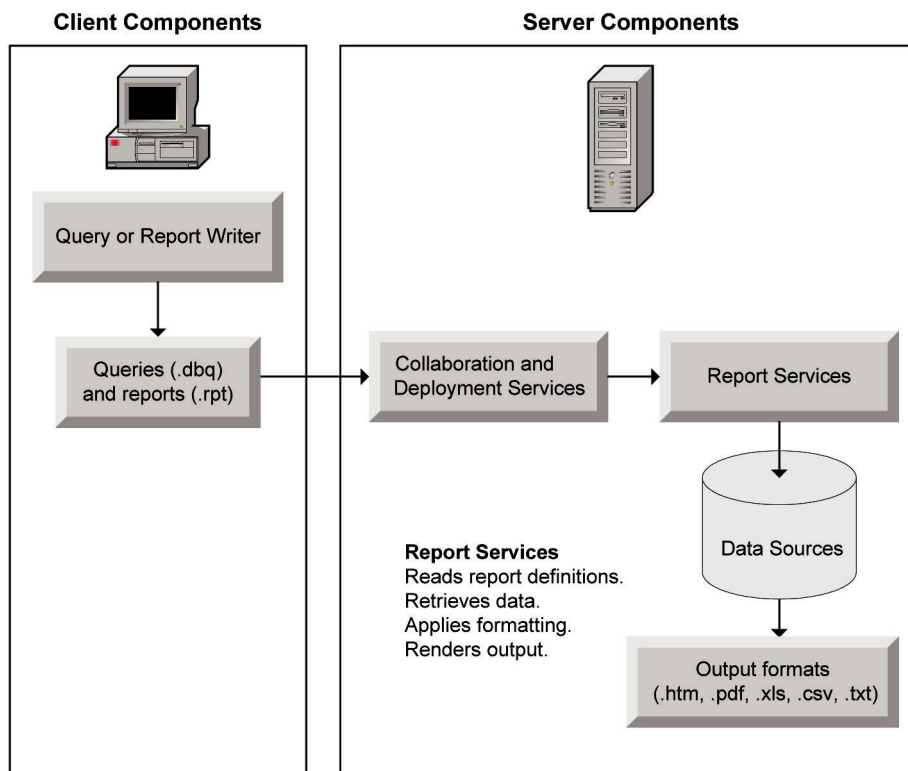
For limitations, see "Compatibility and Limitations" on page 79.

Process

Report Services reads the report, retrieves the data from the data source, performs any data manipulation (such as calculations), and formats the output into a specified format (such as HTML).

Figure 6-1

Report Services process example



Requirements

By default, Report Services is enabled through IBM Collaboration and Deployment Services. All dynamic web and scheduled queries are run via Report Services provided they meet the following criteria:

- Reports cannot use the system variable `&&&SCServerLibrary`.
- Reports cannot contain macros, derived fields, or crosstabs.
- Reports cannot contain the summary operations `Cumulative Sum`, `First`, or `Last`.
- Reports cannot be form or label type reports.
- Reports cannot contain objects inserted with Showcase Report Writer's drawing tool.
- Reports cannot contain images that are used as the report background.
- Reports cannot contain columns with the `Display as Picture` option enabled.
- The output must be a format Report Services supports. Currently supported formats include HTML, Microsoft Excel, PDF, CSV, and delimited text.
- Scheduled queries or reports cannot have the `Send to Default Printer` option selected.

Reports meeting the criteria are processed on the server by Report Services. Reports not meeting the criteria are processed on the Enterprise Server.

Localization Settings

Report Services localizes its output by using the locale set when the JVM initializes.

The Java locale setting is important for language specific behavior such as date and time formats, decimal separators, and so on. In Java, the default locale setting is based on the Java `user.language` and `user.region` system properties. On the IBM i, the `user.language` and `user.region` system properties are based on the language (QLANGID) and country IDs (QCNTRYID) of the IBM i job.

Image and Chart Support

Report Services supports Showcase Report Writer reports (*.rpt*) containing embedded images and supported chart types. The following sections detail embedded image and chart support and identify limitations.

Report Services support for embedded images in reports

- Report files must be in Report Writer 7.0 or later format. Embedded report images created with Report Writer versions prior to 7.0 are not supported.
- Reports cannot contain images inserted with Report Writer's drawing tool.
- Reports cannot contain images that are used as the report background.
- Portable document format (PDF) and Microsoft archived web page (MHT) are the only Report Services output formats that support images. The HTML output does not currently support images, therefore Report Services will generate HTML reports without graphics when the IBM Collaboration and Deployment Services Deployment Portal default output format option is set to HTML.

- The Enterprise Server does not support the HTML or MHT output formats. When HTML or MHT is selected as the output format, the report must be rendered in Deployment Portal using Report Services.
- Internet Explorer can open/render MHT files. Mozilla Firefox requires a plug-in, which is freely available on the Mozilla web site.

***Note:** Report Services does not limit the size of images that are embedded in reports. However, image size does affect performance. An image size of 100 KB or less is recommended.*

Report Services support for charts in reports

The following guidelines define Report Service's support for charts in reports.

- To use Report Services to render Report Writer reports containing charts, Portable App Solutions Environment (PASE) is required on the IBM i. Without PASE, reports containing charts will be rendered with the Enterprise Server.
- Report files must be in Report Writer 7.0 or later format. Charts created with Report Writer versions prior to 7.0 are not supported.
- Reports need to be manually run and then saved to Showcase Warehouse Manager Server with Report Writer 7.0 or later (or to the IBM Collaboration and Deployment Services Repository if using Report Writer 8.0 or 9.0), in order for reports to accurately capture custom chart formatting (custom colors, edit masks, and so on). If a report with a chart, containing custom chart formatting, is not first run in Report Writer 7.0 or later, Report Services may render the report chart with default formatting (for example, a custom edit mask of #, ###.00; #, ###.00 – may default to 0.00). The custom formatting will not display properly until the original report is first run in Report Writer 7.0 or later and then re-saved to Warehouse Manager Server (or to the repository if using 8.0 or 9.0).
- Portable document format (PDF), Microsoft archived web page (MHT), and HTML are the only Report Services output formats that support charts.
- Report Writer does not support the HTML format. A schedule that uses the HTML format will not run successfully on an Enterprise Server.
- Charts based on crosstabs are not supported (they are automatically passed to an Enterprise Server).
- Report Writer reports use special technology for charts. This technology is only compatible with Internet Explorer.
- Certain chart types and options are not supported (they are automatically passed to an Enterprise Server).

The following table defines support for images and charts.

Note: Note: The table below ignores particular chart restrictions and assumes the charts consist of supported types.

| | PDF | HTML | MHT | XLS, TXT, CSV |
|--------------------------|--------------------------------------|--|--|---|
| Tabular only | Report Services or Enterprise Server | Report Services | Report Services | Report Services or Enterprise Server |
| Charts | Report Services or Enterprise Server | Report Services | Report Services | Report Services or Enterprise Server data only (charts do not display) |
| Images | Report Services or Enterprise Server | Report Services | Report Services | Report Services or Enterprise Server data only (images do not display) |
| Charts and images | Report Services or Enterprise Server | Report Services | Report Services | Report Services or Enterprise Server (charts and images do not display) |
| Crosstabs | Enterprise Server | Report Services (crosstabs do not display) | Report Services (crosstabs do not display) | Report Services or Enterprise Server (crosstabs do not display) |

In general, PDF format provides the best support for both images and charts and is also supported by Report Writer. HTML and MHT are not supported by Report Writer; these formats fail if the report file contains features that are not supported by Report Services.

A chart using an unsupported type automatically rolls over to an Enterprise Server if and only if PDF format is selected. Other formats result in a valid tabular result without the chart.

Data Sources

Named connections have been replaced with IBM Collaboration and Deployment Services Deployment Manager data source definitions. See the "Resource Definitions" chapter of the Deployment Manager User's Guide for information on configuring and working with data source definitions.

Default Format Template Definition

The HTML generated by Report Services is clean and easy to read. This improved look is achieved through the default format template. The default format template can be customized to fit individual needs and preferences. This section defines formatting templates and how Report Services uses them.

Overview

Formatting templates are designed to improve the appearance of rendered HTML based on a query (.dbq) definition. A .dbq does not contain enough information to display different borders or backgrounds. Each .dbq is also hard-coded to turn on HTML table grid lines. Templates provide a means to override the default formatting with visually appealing styles.

A goal of format templates is to provide an alternative mechanism for formatting reports. Instead of requiring a .dbq containing formatting, the user can specify a template. The template applies formatting to the various report sections.

Configuration

The default .dbq format template loads from the Default.xml file located in the following directory:

```
<repository_install_dir>\components\reporting-engine\ShowCase\proddata\reporting\
rfotemplate\FormatTemplate
```

The default format template only applies to queries run via Report Services. The template is not applied to Showcase Report Writer reports. When a report is rendered, the formatting template is applied to the report definition to modify the default formatting.

XML Template Files

Formatting templates are stored in XML format. The supported XML attributes for each section are defined later in this section. The default formatting template consists of the following XML elements:

A user may select a different default template at any time by renaming an alternative template to Default.xml. The server does not need to be restarted for the change to take effect.

```
<?xml version="1.0"?>
<FormatTemplate>
  <Table attributes />
  <ColumnHeadings attributes />
  <DetailLines attributes />
  <SeparatorRows attributes />
  <BreakGroups>
    <BreakGroup>
      <Footer attributes />
    </BreakGroup>
    <BreakGroup>
      <Footer attributes />
    </BreakGroup>
  </BreakGroups>
</FormatTemplate>
```

Tag Definitions

Table: Any recognized attribute can be specified. The attributes are mapped directly to the Table HTML tag, or the associated style sheet (for example, setting border or background attributes). The attributes are applied to the entire table containing the output. Typical settings include:

```
TABLE-BORDER="0" CELLSPACING="0".
```

ColumnHeadings: Defines the border and background used for the column headers. Each of the settings for top, bottom, left and right borders are applied to the column headings as a group. Since Showcase Query only has one heading row, it effectively sets the top/bottom of the first and only header row.

DetailLines: Define the attributes applied to each detail cell in the report.

SeparatorRows: Defines the formatting applied to blank separator rows appearing between summary values and the start of the next data set.

BreakGroups: A collection of BreakGroup elements.

BreakGroup: Element containing formatting information related to this break group. The first break group is always the entire report. Each subsequent break group represents the next break level. If the report contains more break levels than what is listed in the template, the inner break levels do not have any additional formatting applied.

Footer: Each footer includes associated attributes. The `BORDER-lrtb-xxx`, `PADDING-lrtb` (where `lrtb` is `LEFT`, `RIGHT`, `TOP`, `BOTTOM`), and `BACKGROUND-COLOR` attributes are supported. Each attribute is applied to the top, bottom, left, or right edge of the group footer. The template controls what border and padding appears around the footer.

Attribute Details

Attribute names are case sensitive, and should be specified in upper case. For more information about each attribute, consult the CSS2 documentation at <http://www.w3.org/TR/REC-CSS2/>.

Report Services supports the following attributes.

"BACKGROUND-COLOR"
"BORDER "BORDER-BOTTOM"
"BORDER-BOTTOM-COLOR"
"BORDER-BOTTOM-STYLE"
"BORDER-BOTTOM-WIDTH"
"BORDER-COLOR"
"BORDER-LEFT"
"BORDER-LEFT-COLOR"
"BORDER-LEFT-STYLE"
"BORDER-LEFT-WIDTH"
"BORDER-RIGHT"
"BORDER-RIGHT-COLOR"
"BORDER-RIGHT-STYLE"
"BORDER-RIGHT-WIDTH"
"BORDER-STYLE"
"BORDER-TOP"
"BORDER-TOP-COLOR"
"BORDER-TOP-STYLE"
"BORDER-TOP-WIDTH"
"BORDER-WIDTH"
"CELLSPACING"
"TABLE-BORDER"

Sample Formats

This section provides sample formatting templates and illustrations of how each template will typically display.

Sample 1 - Simple Borders and Shading

```
<?xml version="1.0"?>
<FormatTemplate>
  <Table TABLE-BORDER="0" CELLSPACING="0" />
  <ColumnHeadings VERTICAL-ALIGN="middle" BACKGROUND-COLOR="gray"
    BORDER-BOTTOM="2 SOLID BLACK" />
  <DetailLines MARGIN="0" />
  <SeparatorRows />
  <BreakGroups>
    <BreakGroup>
      <Footer BORDER-TOP="2 solid black" />
    </BreakGroup>
  </BreakGroup>
```

```

        <Footer BORDER="1 solid black" />
    </BreakGroup>
</BreakGroups>
</FormatTemplate>

```

Figure 6-2

Simple borders and shading sample

| Packagewless | | Channel | Qtr1 Actual | Actual % of budget | Qtr2 Colas |
|---|-----------|---------|-------------|--------------------|------------|
| East | Retail | | **\$12,975 | 100. | 25102. |
| East | Wholesale | | **\$673 | 100. | 2275. |
| Count: 2. Maximum: Wholesale Minimum: Retail | | | | | |
| East | Retail | | **\$4,717 | 100. | 9340. |
| East | Wholesale | | **\$150 | 100. | 506. |
| Count: 2. Maximum: Wholesale Minimum: Retail | | | | | |
| Central | Retail | | **\$9,444 | 100. | 17407. |
| Central | Wholesale | | **\$514 | 100. | 1672. |
| Count: 2. Maximum: Wholesale Minimum: Wholesale | | | | | |

Sample 2 - Complex Borders and Shading

```

<?xml version="1.0"?>
<FormatTemplate>
    <Table TABLE-BORDER="0" CELLSPACING="0" BORDER="2 solid green" BORDER-LEFT="4
solid purple" junk="junkvalue" />
    <ColumnHeadings VERTICAL-ALIGN="middle" BORDER="1 dashed orange"
BORDER-LEFT="3 SOLID PINK" BORDER-RIGHT="3 SOLID blue" />
    <DetailLines MARGIN="0" BACKGROUND-COLOR="gray" BORDER="1 dashed blue"
BORDER-TOP="3 SOLID PINK" BORDER-BOTTOM="3 SOLID green" BORDER-LEFT="3 SOLID
blue" BORDER-RIGHT="3 SOLID red" />
    <SeparatorRows MARGIN="0" BACKGROUND-COLOR="silver" BORDER-TOP="3 dashed
red" BORDER-BOTTOM="3 dashed purple" BORDER-LEFT="3 dashed brown" BORDER-
RIGHT="3 dashed orange"/>
    <BreakGroups>
        <BreakGroup>
            <Footer BORDER="4 solid green" BACKGROUND-COLOR="pink" />
        </BreakGroup>
        <BreakGroup>
            <Footer BORDER="4 solid yellow" BACKGROUND-COLOR="yellow" />
        </BreakGroup>
    </BreakGroups>
</FormatTemplate>

```

Figure 6-3
Complex borders and shading sample

| | Packagename | Channel | Qtr1 Actual | Actual % of budget | Qtr2 Colas |
|---------|-------------|--|-------------|--------------------|------------|
| East | Bottle | Retail | **\$12,975 | 100. | 25102. |
| East | Bottle | Wholesale | **\$673 | 100. | 2275. |
| | Bottle | Count: 2. Maximum: Wholesale Minimum: Retail | | | |
| East | No Pack | Retail | **\$4,717 | 100. | 9340. |
| East | No Pack | Wholesale | **\$150 | 100. | 506. |
| | No Pack | Count: 2. Maximum: Wholesale Minimum: Retail | | | |
| Central | Bottle | Retail | **\$9,444 | 100. | 17407. |
| South | Wholesale | | **\$14 | 100. | 1679. |

Sample 3 - Classic45.xml

This example is the same as deleting the Default.xml file. The formatting is designed to match the default formatting provided in Showcase Query.

```
<?xml version="1.0"?>
<FormatTemplate>
</FormatTemplate>
```

Figure 6-4
Classic sample

| | Packagename | Channel | Qtr1 Actual | Actual % of budget | Qtr2 Colas |
|---------|-------------|--|-------------|--------------------|------------|
| East | Bottle | Retail | **\$12,975 | 100. | 25102. |
| East | Bottle | Wholesale | **\$673 | 100. | 2275. |
| | Bottle | Count: 2. Maximum: Wholesale Minimum: Retail | | | |
| East | No Pack | Retail | **\$4,717 | 100. | 9340. |
| East | No Pack | Wholesale | **\$150 | 100. | 506. |
| | No Pack | Count: 2. Maximum: Wholesale Minimum: Retail | | | |
| Central | Bottle | Retail | **\$9,444 | 100. | 17407. |
| South | Wholesale | | **\$14 | 100. | 1679. |

Compatibility and Limitations

The following sections outline current Report Services limitations:

Edit masks

- Decimal values may not display correctly in IBM Collaboration and Deployment Services Deployment Portal output when run through Report Services. When a Showcase Report Writer report has a numeric column with `Default` set for the edit mask, the edit mask information is not completely saved in the resulting XML file (there is an `EDITMASK` element but no `EDITMASK_STRING` element). As a result, Report Services defaults the edit mask to the locale default of `0.#####`. This causes trailing zero decimal positions to not display (decimal positions do not line up). A work around for this problem is to change the edit mask setting in Report Writer from `Default` to `0.000`
- Character edit masks are not applied.
- The `*` edit mask is not applied. The `*` means repeat the next character to pad the field width.
- For edit masks that change based on locale settings, the locale settings are server based. Modify the edit mask to override for a specific format.
- Exponential notation for edit masks displays three exponent digits instead of the number specified in the edit mask.
- Color edit masks are not supported. For example, in Report Writer the user can choose the edit mask `#,##0.00_);[Red](#,##0.00)` for numeric columns. This edit mask causes negative values to be wrapped in parentheses and display in red. When rendered through Report Services, this edit mask will still wrap negative values in parentheses, but will not display in red.

Format

- The default format template only applies to queries run via Report Services. It will not be applied to Report Writer reports, and is not supported when queries are run via interactive Showcase Query or via the Enterprise Server.
- Report Services writes null values as empty cells for Excel output. This is consistent with how Access, Excel ODBC, and SQL Server transformer handle empty cells.
- For queries, column widths are set based on the font used for the data instead of the font used by the column heading.
- Report Services rounds to the nearest second when nanoseconds are displayed, whereas Query/Report Writer truncate the nanoseconds.
- Use of large borders may result in data truncation in PDF output.
- Report Services does not support the double, standoff, rounded, or shadow border options (a plain single border is substituted instead).
- Detail data located outside of the active print area in Report Writer will render in Report Services HTML output. Users often move chart data off the page when they only want the chart to display in the report (not the data off to the right). There are a few methods that can be applied to ensure the detail text off the page does not render:
 - * Change the off page text to white. This prevents the text from displaying.
 - * Use the `Format/Sections` menu command and hide the detail section.
 - * Create a chart then select and delete the detail data from the screen. The chart will remain valid.

- Objects with left-edge coordinates residing before the left margin are not supported. For example, if a text box is inserted into a header, and the box's left edge is placed before the header's left margin, the text in the text box will not render.
- Report Services does not support white space between sections. In Report Writer, sections can be pulled apart. For example, the user can pull the page footer down from the detail section (possibly for readability purposes). Report Services does not render this separation.
- Report Services does not support overlapping sections. In Report Writer you can overlap sections for effect, or to conserve space. Report Services will not render overlapped sections.

IBM i specific

- On the IBM i, Infinium is not supported via Report Services.
- Japanese characters will not properly display in Adobe PDF documents unless the IBM i server has the Japanese language pack installed, and the report (DBQ/RPT) author explicitly specifies Japanese fonts when creating reports.

On the IBM i, the following font directories are searched:

```
/QIBM/ProdData/OS400/Fonts/TTFonts
/QIBM/ProdData/Java400/fonts/TrueType
/QIBM/UserData/OS400/Fonts/TTFonts
```

Report components

- Report Services does not support reports containing macros, derived fields, or crosstabs.
Note: Although Report Services can often render reports containing unsupported elements (crosstabs for example), the unsupported elements will be absent from the reports (a report containing a crosstab will render, sans the crosstab).
- Report Services does not support queries or reports that reference members. A workaround is to create an alias for the member in Showcase Warehouse Manager Client, then make the query or report reference the alias instead of the member.
- Report Services does not support reports that contain objects inserted with Report Writer's drawing tool.
- Report Services does not support reports that use images as the report background (wallpaper).
- Report Services does not support reports with columns that have the `Display as Picture` option enabled.
- Report Services renders Report Writer images in their original dimensions. Report Writer allows stretching and skewing of images; when rendered in Report Services, images always retain their original dimensions.
- Report Services does not support reports containing the summary operations `Cumulative Sum`, `First`, or `Last`.
- Report Services does not support form or label type reports.
- Queries and reports containing the following functions are processed through the Enterprise Server: `ASCII`, `BIT_LENGTH`, `CONVERT`, `DATABASE`, `REPEAT`, `REPLACE`, `TIMESTAMPADD`.
- After deleting a crosstab from a report and saving to the IBM Collaboration and Deployment Services Repository (Report Services does not support crosstabs), Deployment Portal will still try to process the report through an Enterprise Server rather than Report Services. The report must be recreated without crosstabs in order to run through Report Services.

Variables

- Report Services ignores any variable not defined in the SQL select statement, as allowing variables that are not referenced in the SQL can cause problems.
- Global variables are treated the same as local variables.

Appendix A - Commands

This appendix lists many common IBM i commands, and Showcase Warehouse Manager Server commands for maintaining your IBM i software.

IBM i Commands

Table A-1 lists common IBM i commands for maintaining server libraries.

Table A-1

Common IBM i commands

| Command | Description | Usage |
|------------|---|---|
| ADDLIB | Add a library to your library list. | To add a Showcase Warehouse Manager Server library to the library list when you need to run a server command from the IBM i command line. |
| CFGTCP | Configure TCP/IP. | To configure IBM i TCP/IP. |
| WRKPTFGRP | 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 <code>DSPSYSVAL SYSVAL (QMODEL)</code> ; to find the serial number, use <code>QSRLNBR</code> for the system value. |
| DSPUSRPRF | Display user profile. | To display a user's profile. |
| EDTF | Edit file. | To display or edit IBM i files. |
| WRKACTJOB | View active jobs. | To see if the Warehouse Manager Servers are running or to view active jobs, such as Warehouse Manager Server jobs. |
| WRKOBJLCK | Check for active jobs against a Warehouse Manager Server or Showcase Ess-base Server library. | To check for IBM i jobs that may be active against a Warehouse Manager Server library; in particular, to check for TCP/IP jobs. Useful when installing over the top of an existing Warehouse Manager Server 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. |

Showcase Warehouse Manager Server Commands

The following tables list common Showcase Warehouse Manager Server commands for maintaining a library. For related security information, see Command Authorities on page 67 in Chapter 5.

Note: Before using these commands, add the library to your IBM i library list and ensure that only one Showcase library is in your library list. Run the command `ADDLIB SCSERVER`, where `SCSERVER` is the name of the library running the server you want to work with.

Table A-2*Common Save/Restore commands*

| Command | Description | Usage |
|------------|---|---|
| SAVSRVRINF | Save server information. | To back up the server or move data stored on the server to another library on the same IBM i or a different IBM i. |
| RSTSRVRINF | Restore server information. | To restore the server data saved using the SAVSRVRINF command. |
| VFYMIG | Perform premigration verifications. | To verify the contents of an 8.x save file before performing a restore to IBM Collaboration and Deployment Services. This command generates a set of HTML files in the IBM Collaboration and Deployment Services Repository that lists the expected results of the restore. |
| MIGPESWIN | Restore Enterprise Reporting for Windows save file. | To restore an Enterprise Reporting for Windows save file into an IBM i installation of IBM Collaboration and Deployment Services. |
| RELODSCCFG | Reload server configuration. | To reload the server configuration without stopping and restarting the server. |

Table A-3*Common TCP/IP and firewall commands*

| Command | Description | Usage |
|-----------|---|--|
| CHGSCSVRA | Add TCP/IP connectivity support. | To enable or disable TCP/IP support after installation. Requires *ALLOBJ, *SECADM, *IOSYSCFG, *JOBCTL, and *SAVSYS authorities. Note: The CHGSCSVRA command will automatically stop the server, disconnect all current users, and stop current jobs. It will also restart the server after changes have been made. |
| ENDSCSVR | End the server. | To end TCP/IP. For example, if you need to restart TCP/IP, you should end TCP/IP first. Requires *ALLOBJ, *SECADM, *IOSYSCFG, *JOBCTL, and *SAVSYS authorities. |
| ADDSCTCP | Configure TCP/IP settings for the Warehouse Manager Server. | To configure the subsystem, TCP/IP ports, and the RMI host name used for the Warehouse Manager Server. Configurable ports include the main listening port, the firewall listening port, and the compressed data port. The ports are added to the IBM i service table and can be viewed with the WRKSRVTBLE command by looking for entries beginning with SCSERVER, SCFIREWALL, and SCRMICOMP and ending with the name of the Warehouse Manager Server library. The RMI host name is stored in an Showcase internal object. Note: Use RMVSCTCP to remove existing configuration settings first. |

| Command | Description | Usage |
|-----------|---|--|
| RMVSCTCP | Remove TCP/IP settings for the Warehouse Manager Server. | To remove the subsystem entries, TCP/IP ports, and the RMI host name used for the Warehouse Manager Server. Note: Administrators should manually remove any entries (if present) in the <i>/<SCSERVER>/Prod-Data/config/showcase.cfg</i> file (where <i>SCSERVER</i> is the name of the Warehouse Manager Server installation library) to ensure that any pre-existing firewall configuration is completely removed. |
| RSTSMPLE | Restore sample databases. | To install the sample databases if not previously installed, or to restore the sample databases back to their original structure (if database files were altered or destroyed). |
| SETJDEOWA | Configure the server on the IBM i for OneWorld integration. | To configure the server on the IBM i for OneWorld integration. The command is also used to enable/disable JD Edwards OneWorld connections for users who do not have an IBM i user profile. For more information, see the topic Non-IBM i Profile Sign-On in Chapter 4 on page 36. |
| STRSCSVR | Start the server. | After an IPL or ENDSCSVR. Requires at least *JOBCTL authority. |
| STRTCP | Start IBM i TCP. | To start IBM i TCP/IP. IBM i TCP/IP must be started before you can enable or start TCP/IP support. |

Table A-4

Common information commands

| Command | Description | Usage |
|------------|---|---|
| CHKSRVRCAT | Check server catalog information. | To inspect Warehouse Manager Server catalogs and delete obsolete Warehouse Manager Server settings. |
| CHGSRVRSTS | Change server status. | To change the status of the server library (for example, <i>SCSERVER</i>) to and from a restricted state. A restricted state prevents users from connecting to the server library. |
| DSPSCVER | Display Warehouse Manager Server version. | To see the version of Warehouse Manager Server installed and other server information. |
| DSPSCINF | Display Showcase hints and tips. | To locate hints and tips provided with Showcase. |

Table A-5

Common options commands

| Command | Description | Usage |
|------------|-------------------------------|---|
| CHGSRVRDFT | Change server defaults. | To change the default *PUBLIC authority to all libraries or tables and enforce Warehouse Manager Server security rules on *ALLOBJ profiles. |
| CHGUSRQRYA | Change user query attributes. | Sets the QAQQINI library name for a given user, group, or the *PUBLIC group. Provides options for specifying different QAQQINI files for each SCTCPSRV job. |

| Command | Description | Usage |
|-----------|---|---|
| SETLIBACC | Set library list access. | To restrict users so they can access only the libraries in their IBM i library lists, or to give users access to all libraries on the IBM i to which they are authorized. |
| SCCHKPWD | Manage IBM I passwords stored in control tables (see SCCHKPWD Command Parameters on page 87). | Update the Showcase 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 IBM I command line, and press F1 for parameter information. |

Table A-6

Showcase Warehouse Builder commands

| Command | Description | Usage |
|---------|--|--|
| MAPDIST | Update Showcase Warehouse Builder definitions. | Updates Warehouse Builder definitions by mapping the source or target data source to a new data source. The utility was created to resolve RSTSRVRINF limitations. |
| 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. |

Table A-7

Server patch rollback command

| Command | Description | Usage |
|----------|---|---|
| RMVSCPTF | Roll back Warehouse Manager Server patches. | To roll back (unapply) Warehouse Manager Server patches. Contains no parameters. Patches can be rolled back in succession to the installation release level. Note: The patch rollbacks rely on updated patch information in the SCPTFHST SQL table and the availability of the backup save files and user space (created when patches are applied) in the Warehouse Manager Server installation library. If any of these is not found when the patch rollback is run, the rollback will not continue. |

Table A-8*Uninstall command*

| Command | Description | Usage |
|----------|--|--|
| UNINSTSC | Uninstall an Warehouse Manager Server library. | To remove a library from the system. For more information about the UNINSTSC save file, see Uninstalling a Server Library on page 51 in Chapter 5. |

Table A-9*Install adapter command*

| Command | Description | Usage |
|----------|--|--|
| DPLSCPKG | Deploy the IBM Collaboration and Deployment Services - Server Adapter for Showcase on the IBM i. | The Deploy Showcase Package (DPLSCPKG) command launches the installer for the Server Adapter for Showcase on the IBM i for use with the IBM Collaboration and Deployment Services Server running in WebSphere. |

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 Showcase Warehouse Builder distributions.

Note: Most Showcase applications support 128-character, mixed-case passwords. JD Edwards One-World supports 10-character, uppercase passwords. For complete information on IBM i password guidelines and restrictions, see your IBM i documentation.

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-10 lists the optional SCCHKPWD parameters.

Table A-10*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 affected RDBs, servers, and data sources is included. |

| Parameter | Description |
|---------------------------|--|
| 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 items using data sources to the specified RDB. (Non-Showcase data sources must use the ODBC parameter to update passwords.) Special values supported are: *LOCAL, which references the local IBM i 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 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 a data source for which passwords may need to be checked or updated. This will affect both scheduled items using DBQs that reference non-Showcase data sources and data sources used in Warehouse Builder 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-Showcase) ODBC-related passwords should be checked or updated. |
| ESSBASE (default = *NONE) | This parameter is no longer supported. |
| 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. |

Appendix B - Sample Databases

This appendix describes the sample databases you can install with the server on the IBM i. The databases are:

- `SCSAMPLE90`, a more summarized database in a star-schema format that represents an OLAP-type database. This database is used with the Showcase product tutorials.
- `SCDBTRAIN`, a fairly normalized database that represents an OLTP-type database, used for training sessions.
- `SCDB1`, a database used for Technical Support purposes.

Installing the Sample Databases

When you install the server on the IBM i, you have the option to install the sample databases. Selecting this option creates one or more libraries on the IBM i, depending on how many of the sample databases you install. For example, choosing to install `SCSAMPLE90` creates a library called `SCSAMPLE90`.

The files (tables) for the sample databases are installed in the corresponding sample libraries. To install one or more of the sample libraries on the IBM i, use one of the following methods:

- Select the option to install the sample libraries as part of the initial installation of the server.
- Run the Restore Sample Library (`RSTSMPLB`) command on the IBM i command line. For more information, see “Installing the Sample Library after Installation with `RSTSMPLB`” on page 97.

Contents of the `SCSAMPLE90` Sample Database

The `SCSAMPLE90` sample database is the basis for the Showcase tutorials. The database contains data for a fictitious company called The Sample Outdoor Company. The Sample Outdoor Company 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 Sample Outdoor Company 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 you can install with the server. The sample database tables are installed into the `SCSAMPLExx` library, where `xx` is the current product version (for example, `SCSAMPLE90`).

The following figure depicts the general structure of the sample database:

Figure B-1

Sample database structure



SCSAMPLE90 Database

ORDERS Table

The ORDERS table is the basis of all transaction information. It contains the order number, location placing the order, date of the order, promised delivery date, date shipped, and so on. 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 originated.

The ORDERS table also contains product-related specifics of each order, such as the quantity ordered, quantity shipped, quantity backlogged, and quantity canceled. 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 Sample Outdoor Company 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 US 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 CA 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 | Char(8) | Third-tier product category |
| QTYORD | Quantity Ordered | Quantity Ordered | Decimal(15) | Quantity ordered |
| QTYSHIP | Quantity Shipped | Quantity Shipped | Decimal(15) | Quantity shipped |
| QTYBACK | Quantity Backordered | Quantity Backordered | Decimal(15) | Quantity backordered |
| QTYCAN | Quantity Cancelled | Quantity Cancelled | Decimal(15) | Quantity cancelled |
| 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 regarding The Sample Outdoor Company's clients. It includes names and addresses and divides the clients into various geographical categories. The Sample Outdoor Door Company consists of three 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 three 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 |

| Column | Description | Heading | Data type | Contents |
|------------|------------------------|----------------------|-----------|-------------------------------|
| 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 Sample Outdoor Company's product information and groups products hierarchically. Each product has three associated levels. The first identifies a product group (for example, clothing, biking, skiing, etc.), the second identifies a type (for example, outerwear, shirts, pants, etc.), and the third identifies 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 Level 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 Level 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 Level 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 Level 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 Level 5 | Varchar(27) | Fifth tier product description |

| Column | Description | Heading | Data type | Contents |
|-------------------|-----------------|-----------------|---------------|--|
| 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 three divisions, 12 regions, 70 areas, and three 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, enabling them to 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 |

| Column | Description | Heading | Data type | Contents |
|------------|-------------|------------|-------------|---------------------------|
| 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) |

DATETYPES Table

The DATETYPES table is used to demonstrate how the product applications handle date values stored as a data type other than DATE. You can use this table independently or join it to another table, or any column, that is defined as a DATE data type. Date values include all dates from 1/1/1998 through 12/31/2008.

This table contains 4,018 rows of data.

Table B-6

DATETYPES table

| Column | Description | Heading | Data type |
|------------|-------------------------------------|-----------------|------------|
| DATEDATE | Date Date | Date Date | Date |
| TSTAMP | Timestamp | Timestamp | Timestamp |
| YY | Two Digit Year Value | YY | Decimal(2) |
| YY_C | Two Character Year Value | YY (char) | Char(2) |
| MM | Two Digit Month Value | MM | Decimal(2) |
| MM_C | Two Character Month Value | MM (char) | Char(2) |
| DD | Two Digit Day Value | DD | Decimal(2) |
| DD_C | Two Character Day Value | DD (char) | Char(2) |
| YYYY | Four Digit Year | YYYY | Decimal(4) |
| YYYY_C | Four Character Year | YYYY (char) | Char(4) |
| EXCELSRL | Five Digit Excel Serial Date | EXCELSRL | Decimal(5) |
| HYF | Five Digit Hundred Year Format Date | HYF | Decimal(5) |
| YYDDD | Five Digit YYDDD Date | YYDDD | Decimal(5) |
| YYDDD_C | Five Character YYDDD Date | YYDDD (char) | Char(5) |
| YYMMDD | Six Digit YYMMDD Date | YYMMDD | Decimal(6) |
| YYMMDD_C | Six Character YYMMDD Date | YYMMDD (char) | Char(6) |
| YYDDMM | Six Digit YYDDMM Date | YYDDMM | Decimal(6) |
| YYDDMM_C | Six Character YYDDMM Date | YYDDMM (char) | Char(6) |
| MMDDYY | Six Digit MMDDYY Date | MMDDYY | Decimal(6) |
| MMDDYY_C | Six Character MMDDYY Date | MMDDYY (char) | Char(6) |
| DDMMYY | Six Digit DDMMYY Date | DDMMYY | Decimal(6) |
| DDMMYY_C | Six Character DDMMYY Date | DDMMYY (char) | Char(6) |
| YYYYDDD | Seven Digit YYYYDDD Date | YYYYDDD | Decimal(7) |
| YYYYMMDD | Eight Digit YYYYMMDD Date | YYYYMMDD | Decimal(8) |
| YYYYMMDD_C | Eight Character YYYYMMDD Date | YYYYMMDD (char) | Char(8) |
| YYYYDDMM | Eight Digit YYYYDDMM Date | YYYYDDMM | Decimal(8) |
| YYYYDDMM_C | Eight Character YYYYDDMM Date | YYYYDDMM (char) | Char(8) |
| MMDDYYYY | Eight Digit MMDDYYYY Date | MMDDYYYY | Decimal(8) |
| MMDDYYYY_C | Eight Character MMDDYYYY Date | MMDDYYYY (char) | Char(8) |
| DDMMYYYY | Eight Digit DDMMYYYY Date | DDMMYYYY | Decimal(8) |

| Column | Description | Heading | Data type |
|------------|-------------------------------|-----------------|------------|
| DDMMYYYY_C | Eight Character DDMMYYYY Date | DDMMYYYY (char) | Char(8) |
| CYYDDD | Six Digit CYYDDD Date | CYYDDD | Decimal(6) |
| CYYMMDD | Seven Digit CYYMMDD Date | CYYMMDD | Decimal(7) |
| CYYDDMM | Seven Digit CYYDDMM Date | CYYDDMM | Decimal(7) |

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 4,018 rows of data.

Table B-7

TIMEDIM table

| Column | Description | Heading | Data type | Contents |
|------------|-------------|------------|-------------|--------------------|
| TDATE | DATE | TDATE | Date | Date |
| TYEAR | YEAR | TYEAR | Integer | Year |
| TMONTH | MONTH | TMONTH | Integer | Month number |
| TDAY | DAY | TDAY | Integer | Day of month |
| TDAYOFWEEK | DAYOFWEEK | TDAYOFWEEK | Integer | Day of week number |
| TDAYNAME | DAYNAME | TDAYNAME | Varchar(9) | Day |
| TDAYOFYEAR | DAYOFYEAR | TDAYOFYEAR | Integer | Day of year number |
| TWEEK | WEEK | TWEEK | Integer | Week number |
| THOLCODE | HOLCODE | THOLCODE | Integer | Holiday code |
| THOLIDAY | HOLIDAY | THOLIDAY | 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 | Decimal(15,0) | Projected unit sales |
| ACT_UNITS | Actual Unit Sales | Actual Units | Decimal(15,0) | Actual units sold |
| PRJ_SALES | Projected Sales | Projected Sales | Decimal(15,4) | Projected sales amount |
| ACT_SALES | Actual Sales | Actual Sales | Decimal(15,4) | Actual sales amount |
| QUOTA | Quota | Quota | Decimal(15,4) | Sales quota |

EMPLOYEES Table

The EMPLOYEES table contains human resources information regarding The Sample Outdoor Company'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 | Char(100) | Surname of employee |
| FIRSTNAME | Given Name | First Name | Char(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 |

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 `SCSERVER/RSTSMPLB` at the command prompt (where `SCSERVER` is the library name in which the server is installed) and press Enter.

The command has three parameters:

- `Database(s) to restore.` Type the name of the database save file to restore. For example, type `SCSAMPLE90`. To restore more than one sample database at once, type `+` for more values and press Enter. The possible values are `SCSAMPLE90`, `SCDB1`, and `SCDBTRAIN`.
- `Clear Target Lib.` If `*YES` is chosen, the sample database library (if it exists) is cleared (all files within destroyed) before the sample files are restored. Sample files are restored to their original version installed with this release.
- `Preserve Logical Files.` If `*YES` is chosen, all logical files on the system that are dependent on the sample files within the sample database library are preserved. Logical files that are dependent on nonproduct files within the sample library are not preserved. To preserve logical files not dependent on sample files, choose `*NO`.

Appendix C - Troubleshooting

This appendix lists common problems and questions regarding Showcase 9 applications and provides resolutions and answers. For other release notes not discussed here, see Chapter 2 in the *Showcase 9 Installation Guide* or log on to the Technical Support Site at www.helpsystems.com/showcase for the latest release notes.

The following topics are covered:

- Installation Troubleshooting
- Showcase Query Troubleshooting
- Showcase Report Writer Troubleshooting
- Showcase Warehouse Builder Troubleshooting
- Showcase Warehouse Manager Server Troubleshooting

Installation Troubleshooting

The following issues detail common installation problems and solutions.

Failed Installation

Problem: Showcase fails to install. The JOBLLOG contains a message that states two objects did not restore.

```
279 objects restored. 2 not restored to SCSERVER.  
A critical restore operation failed during installation An error  
occurred installing Server version X.  
Warehouse Manager failed to install in to library SCSERVER One or  
more product features failed to install or configure  
The Server and related components failed to install in to library  
SCSERVER
```

Resolution: The IBM i system value called QFRCCVNRST (Force Conversion on Restore) determines whether executables are retranslated during the restore to the system. Change the system value QFRCCVNRST to '2' to allow Showcase to install.

Showcase Query Troubleshooting

The following issues detail common Showcase Query problems and solutions.

Passwords

Problem: A message similar to Password not correct for user profile xxxxxx is received while attempting to connect to pre- 9.0 data source that was created or modified after 9.0 was installed on the client.

Resolution: Modify the data source registry entry to use the drivers for the correct pre- 9.0 version.

Update the path in the following registry key to point to the SCOJDBC.DLL driver in the correct Showcase version's \Bin directory:

```
HKEY_CURRENT_USER\Software\ODBC\ODBC.INI\  
"datasource name"\Driver
```

(for example, C:\Program Files\ShowCase Suite\bin\SCOJDBC.DLL)

vista.SAV

Problem: After installing 9.0 from a network, an error message displays while attempting to run Query:

The save file, "\\networkDir\TestDrv\.....\vista.SAV' could not be opened.

The directory specified in the registry for the location of the save does not exist.

Change the path specified at HKEY_CURRENT_USER\Software\Showcase\Query\OptionsVistaSessionSaveFile to a valid directory.

Resolution: Update the registry to point to an existing vista.SAV file. When performing installations from the network (*after* installing to the network and *before* running an installation on a local PC), update the registry to point to one of the following items:

- The shared vista.SAV file (if variables are to be shared across the network).
- The local installation folder that users will use for their network installation (for example, C:\Program Files\ShowCase\...). If the error is seen after the installation to the local PC, update the registry on the local PC to point to the correct location for the shared global file vista.SAV or remove the entry. Query will create a new file.

Update the path of the VistaSessionSaveFile specified in the registry at

```
HKEY_CURRENT_USER\Software\Showcase\Query\Options.
```

Note: The end user must have read-write permission to the specified location.

Members and Data View Manager

Problem: How can I specify a member using Data View Manager? **Resolution:** This is not possible. IBM i SQL views do not support members.

Workaround: Placing a system alias on a member through the Showcase Warehouse Manager Client interface will allow you to access a file's member in a data view and through Showcase Warehouse Builder. The member with the alias appears as a table in Query and can be used just like any other table.

IBM i SQL does not provide syntax for directly specifying a file member. Therefore, access any file member other than *FIRST by implementing a database override. Note that an override applies to the `_file_name_`, so whenever you specify that filename, the override takes effect. Only one override can be specified at a time. If two overrides are issued, the second one takes precedence over the first. As a result, two different members cannot be joined from the same physical file. Further, you cannot join nonfirst members from two different physical files if the files have the same name (the override is specific to filename only).

To work around this limitation, copy the data from the desired file member into another file. A SQL view will not allow you to access nonfirst file members from views. Try one of the following actions:

- From an IBM i session, use the `CPYF` command to specify from which member to pull. Use the command to also create the new file.
- From the Showcase desktop, run a basic, select-all-fields query against the nonfirst member (without any joins) and send that data into a batch output file. Join between the *FIRST member of the original file and the batch output file, or create multiple batch output files if you need to join multiple nonfirst members. With this option, remember to refresh the batch output file whenever you need to run the query.

Showcase Report Writer Troubleshooting

The following issues detail common Showcase Report Writer problems and solutions.

Graph in Break Group Headers

Problem: Graph in break group header reflects incorrect data. When placing a graph in a break group header, the graph reflects data from the last group on the page (not the data in the group immediately under the header).

Resolution: Inserting a new page after each group footer, to keep each group on its own page, reflects the data in the chart correctly. Report Writer still displays the data in the wrong order, but the chart is accurate. The values are sorted in alphabetical order in the graph, whereas the report data are not sorted because it is a dimension, not a column.

Showcase Warehouse Builder Troubleshooting

The following issues detail common Showcase Warehouse Builder problems and solutions.

Save and Restore Individual Sets/Definitions

Problem: How do I save and restore individual sets/definitions? If I currently have a development and a production control server, and I saved a definition set to the development control server, how would I easily migrate the definition set to the production control server? The `SAVSRVRINF` and `RSTSRVRINF` commands save and restore the entire library.

Resolution: See "Restoring Server Information" on page 54 for details regarding saving/restoring individual definitions and sets.

Showcase Warehouse Manager Server Troubleshooting

The following issues detail common Showcase Warehouse Manager Server problems and solutions.

Passwords

Problem: An error message is received similar to `Password not correct for user profile when connecting to IBM i`.

Resolution: This message is displayed when attempting to connect an 9.0 client to a pre- 9.0 server version, which is not supported.

Journals

Problem: Is there a way to combine the journals in the server library into one journal?

Resolution: Combining of journals is not supported.

Index

A

| | |
|---|--------|
| ADDLIBLE Command | 83 |
| ADDSTCP Command | 50, 84 |
| AOS Parameter | |
| SCCHKPWD Command | 88 |
| Application-Level Security | 9 |
| Applying Column Filters | 10 |
| Applying Row Security | 10 |
| APPSECOPT Parameter | |
| SETJDEOWA Command | 41 |
| Assigning Job-Run Priorities | 14 |
| Assigning Query Run Limits | 13 |
| Assigning the Level of Parallelism | 15 |
| Auditing | |
| Analyzing | 18 |
| Collecting Information | 18 |
| Impact to Overall Performance | 17 |
| Indexes Advised | 19 |
| Most Frequently Queried Files | 20 |
| Query Governor Informatio | 20 |
| SELECT Statements | 18 |
| Auditing Environment | |
| Managing | 20 |
| Auditing to Understand Your Queries | 17 |
| AUTHSYSNAM Parameter | |
| SETJDEOWA Command | 41 |
| AUTHSYSPWD Parameter | |
| SETJDEOWA Command | 41 |
| AUTHSYSUSR Parameter | |
| SETJDEOWA Command | 41 |

C

| | |
|--|----|
| CFGTCP Command | 83 |
| CHGSCSVRA Command | 84 |
| CHGSRVRDFT Command | 85 |
| CHGSRVRSTS Command | 85 |
| CHGUSRQRYA Command | 85 |
| CHKSRVRCAT Command | 85 |
| Command Authorities | 67 |
| Configuring Firewall Settings with ADDSTCP | 50 |
| Controlling Level of Parallelism | 15 |
| Controlling System Usage | 12 |
| Creating Indexes | 23 |

D

| | |
|-------------------------------------|----|
| DB2 Symmetric Multiprocessing | 17 |
|-------------------------------------|----|

| | |
|-------------------------------|----|
| Dependent Data Views | |
| Changing Data Views | 32 |
| Changing Physical Files | 31 |
| DETAILS Parameter | |
| SCCHKPWD Command | 87 |
| DPLSCPKG Command | 87 |
| DRDA | 69 |
| DSPSCINF Command | 85 |
| DSPSCVER Command | 85 |
| DSPSYSVAL Command | 83 |
| DSPUSRPRF Command | 83 |

E

| | |
|---------------------------|----|
| EDTF Command | 83 |
| Enable Expert Cache | 16 |
| ENDSCSVR Command | 84 |
| ESSBASE Parameter | |
| SCCHKPWD Command | 88 |

I

| | |
|---|----|
| IBM i Commands | 83 |
| IBM i Passwords | |
| Changing with an ShowCase Application | 32 |
| Updating with the SCCHKPWD Command | 32 |
| IBM i Security Considerations | 5 |
| IBM i Work Management | 64 |
| Information Workflow | 33 |
| Installing the Sample Database | 89 |

J

| | |
|--|----|
| JD Edwards (JDE) Integration | 36 |
| JDE Security | 36 |
| Multiple Library List Environments | 42 |
| User-Defined Codes (UDC) | 36 |
| JDE Security | 36 |
| Activating Inclusive Row Security | 38 |
| Application Security | 37 |
| Business Unit Security | 39 |
| Column-Level Application Security | 37 |
| Column-Level Security | 38 |
| Inclusive Row Security | 38 |
| Object-Level Security | 37 |
| Role Relationship Support | 40 |
| Row-Level Security | 38 |
| JDEOWPGMLB Parameter | |
| SETJDEOWA Command | 41 |

| | |
|-------------------------|----|
| JDEOWRLS Parameter | |
| SETJDEOWA Command | 40 |
| JDEOWRSLB Parameter | |
| SETJDEOWA Command | 42 |
| JDEOWSYSLB Parameter | |
| SETJDEOWA Command | 41 |
| Job-Run Priorities | |
| Assigning | 14 |
| Setting Server | 14 |
| Journals | 62 |

L

| | |
|----------------------|----|
| Level of Parallelism | |
| Assigning | 15 |
| Controlling | 15 |

M

| | |
|-------------------------------------|----|
| Managing | |
| Performance | 5 |
| Relational Database | 5 |
| Managing Data | 29 |
| Restoring Data | 31 |
| MAPDIST Command | 86 |
| MAPDIST Utility | 57 |
| MIGPESWIN Command | 84 |
| Moving Data from Test to Production | |
| Redirecting Data | 29 |
| Saving and Restoring | 29 |
| Multiple QAQQINI File Support | 35 |

N

| | |
|------------------------|----|
| NEWPWD Parameter | |
| SCCHKPWD Command | 88 |

O

| | |
|------------------------|----|
| ODBC Parameter | |
| SCCHKPWD Command | 88 |

P

| | |
|--------------------------|----|
| Performance Analyzer | |
| analyze your query | 22 |
| tune your query | 22 |
| PWDCONFIRM Parameter | |
| SCCHKPWD Command | 88 |

Q

| | |
|-------------------------|----|
| Query Optimizer | 21 |
| Query Performance | 21 |

R

| | |
|--|--------|
| RDB Parameter | |
| SCCHKPWD Command | 88 |
| Redirecting Data | 29 |
| Reloading the Server Configuration | 52 |
| RELODSCCFG Command | 52, 84 |
| Restricted State | |
| Applying | 59 |
| Removing | 60 |
| Restricting Users | 9 |
| RMVSCPTF Command | 86 |
| RMVSCTCP Command | 85 |
| RSTSMLB Command | 85, 97 |
| RSTSRVRINF Command | 84 |
| Run Limits | |
| Assigning Query | 13 |
| Setting Maximum | 13 |
| Verifying | 14 |
| RUNAUTOVR Parameter | |
| SETJDEOWA Command | 41 |

S

| | |
|---------------------------------------|--------|
| Sample Database | |
| Contents | 89 |
| CUSTOMERS Table | 92 |
| DATETYPES Table | 95 |
| EMPLOYEES Table | 97 |
| Installing | 89 |
| ORDERS Table | 90 |
| PRODUCTS Table | 93 |
| SALESORG Table | 94 |
| SALESPERF Table | 96 |
| SALESREPS Table | 94 |
| TIMEDIM Table | 96 |
| Sample Databases | 89 |
| SAVSRVRINF Command | 84 |
| SBMDD Command | 86 |
| SCCHKPWD Command | 86, 87 |
| Security Considerations | |
| Application Output | 11 |
| Application-Level Security | 9 |
| Applying Column Filters | 10 |
| Applying Row Security | 10 |
| Designating Data Administrators | 7 |
| IBM i | 5 |
| Menu-Level Security | 8 |
| Restricting Users | 9 |
| Server Security | 6 |
| Setting Application Security | 11 |
| User Profile Scheme | 8 |
| Server Configuration | |
| Reloading | 52 |

| | | | |
|--|--------|---|-----|
| Server Information | | Contacting | 1 |
| Restoring | 54 | Website | 1 |
| Saving | 53 | Troubleshooting | |
| Server Library Usage | 50 | Installation | 99 |
| Server Logs | 61 | Showcase Query | 99 |
| Server Patches | | Showcase Report Writer | 101 |
| Unapplying | 61 | Showcase Warehouse Builder | 101 |
| Set Client Session Timeout Limits | 13 | Showcase Warehouse Manager Server | 101 |
| SETJDEOWA Command | 40, 85 | vista.SAV | 100 |
| SETLIBACC Command | 86 | Tuning System Performance | 16 |
| Setting Application Security | 11 | | |
| Setting Maximum Query Run Limits | 13 | U | |
| Setting Server Job-Run Priorities | 14 | Unapplying Server Patches | 61 |
| Showcase Enterprise Server | | UNINSTSC Command | 87 |
| Starting | 4 | UPDATE Parameter | |
| Stopping | 4 | SCCHKPWD Command | 87 |
| Showcase Warehouse Manager Server | | User Exit Programs | 67 |
| Starting | 3 | User-Defined Codes (UDC) | 36 |
| Stopping | 3 | | |
| Showcase Warehouse Manager Server Commands | 83 | V | |
| showcase.cfg | 69 | Verifying the Run Limit | 14 |
| Single Sign-On | 62 | vista.SAV | 100 |
| SINGLUSRPRF Parameter | | | |
| SETJDEOWA Command | 41 | W | |
| Starting Servers | | Warehouse Builder | |
| Showcase Enterprise | 4 | Optimize Performance | 25 |
| Showcase Warehouse Manager | 3 | Work Management Controls | |
| Stopping Servers | | Job Class | 64 |
| Showcase Enterprise | 4 | Job Descriptions | 66 |
| Showcase Warehouse Manager | 3 | Routing Entries | 65 |
| TCP/IP | 51 | Storage Pools | 65 |
| STRDD Command | 86 | Subsystems | 65 |
| STRSCSVR Command | 85 | User Profiles | 66 |
| STRTCP Command | 85 | WRKACTJOB Command | 83 |
| | | WRKOBJLCK Command | 83 |
| T | | WRKPTFGRP Command | 83 |
| TCP/IP | 49 | WRKSRVTBLE Command | 83 |
| Stopping | 51 | | |
| Technical Support | | | |

