

WebFOCUS

Resource Governor Administrator's
and User's Manual
Release 7709

Active Technologies, EDA, EDA/SQL, FIDEL, FOCUS, Information Builders, the Information Builders logo, iWay, iWay Software, Parlay, PC/FOCUS, RStat, Table Talk, Web390, WebFOCUS, WebFOCUS Active Technologies, and WebFOCUS Magnify are registered trademarks, and DataMigrator and Hyperstage are trademarks of Information Builders, Inc.

Adobe, the Adobe logo, Acrobat, Adobe Reader, Flash, Adobe Flash Builder, Flex, and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Due to the nature of this material, this document refers to numerous hardware and software products by their trademarks. In most, if not all cases, these designations are claimed as trademarks or registered trademarks by their respective companies. It is not this publisher's intent to use any of these names generically. The reader is therefore cautioned to investigate all claimed trademark rights before using any of these names other than to refer to the product described.

Copyright © 2019, by Information Builders, Inc. and iWay Software. All rights reserved. Patent Pending. This manual, or parts thereof, may not be reproduced in any form without the written permission of Information Builders, Inc.

Contents

Preface	7
Conventions	8
Related Publications	9
Customer Support	9
Information You Should Have	10
User Feedback	11
iWay Software Training and Professional Services	11
1. What Is Resource Governor?	13
Resource Governor Overview	13
Resource Governor Features	14
Resource Governor Operations	15
How Resource Governor Works	16
Resource Governor Usage Monitoring	16
Setting Rule Parameters With Resource Governor.....	18
Building Rules.....	18
Creating Business Rules.....	18
Reporting With Resource Governor.....	19
2. Configuring and Administering Resource Management	21
Resource Management Interface	21
Ribbon.....	21
Configuring Resource Management	22
Configuring Many-to-One Monitoring.....	30
Configuring Resource Governor.....	31
Administering Resource Management From the Web Console	32
Resource Management Configuration Maintenance.....	35
Enabling or Disabling Monitoring.....	35
Logging.....	35
Archive Log Files.....	38
List Unarchived Log Files.....	47
Working With Repository Maintenance.....	47
Reconfiguring the Repository.....	57

Migration	59
Configuration Maintenance	60
Resource Management Properties	60
Global Governing	70
Configuring Application Paths	70
Migrating Resource Management	70
Working With License Keys	71
Removing Resource Management	71
Working With Application Directories	72
Monitoring at the Application Level	72
Monitoring Individual Data Sources	75
Monitoring Relational Adapters	79
Relational Adapter Properties	79
Monitoring Relational Data Adapter Sources	81
Monitoring at the Data Service Level	86
Migrating the Resource Management Repository	88
Requirements for Migrating the Resource Management Repository	89
Migrating a Resource Management Repository for a New Server Install	89
Migrating a Resource Management Repository After a Server Refresh	94
Migrating a Resource Management Repository on MVS	99
Event Routing	102
3. Rules Files and Knowledge Bases	103
Overview	103
Working With Business Rules	104
Creating Business Rules	104
Editing Business Rules	109
Working with Resource Thresholds	111
Creating Thresholds	111
Editing Resource Thresholds	114
Working with Knowledge Bases	115
Applying a Global Knowledge Base	115
Applying an Individual Knowledge Base	118

Maintaining Knowledge Bases	122
4. Resource Governor Report Options	129
Reports Overview	129
Resource Governor Reporting	130
Resource Governor Report Options	138
Report Filters.....	138
Online Help.....	141
Drilling Down With Reports.....	142
Monitor Preference and Reports.....	142
Alternate Reporting Output.....	142
General Reports	143
Monitored Sessions.....	143
Monitored Commands.....	144
Repository Statistics.....	144
Data Sources Never Used.....	145
Procedures Never Used.....	145
ACI Monitoring.....	146
Resource Governor Reports	147
Monitored Data Sources.....	147
Monitored Procedures.....	148
Resource Thresholds Report.....	148
Knowledgebase Properties Report.....	149
Threshold Advisor Report.....	150
Governed Requests Report.....	150
5. Getting Help	153
Resource Management Administration Help	153
Web Console Help.....	153
Reporting Application Context-Sensitive Help From Reports	155
A. Administrative Usage Monitoring Tables Column Descriptions	157
Usage Monitoring Table (RMLDB.MAS)	157
SMSERVERS Segment.....	157
SMSESS Segment.....	158

SMRPCS Segment.	162
SMRPCCMD Segment.	165
SMRPCWF Segment.	165
SMQUERY Segment.	166
SMCMDLN Segment.	170
SMGOVEND Segment.	170
SMRMTS Segment.	171
SMFNCTNS Segment.	172
SMRELTNS Segment.	173
SMFROMS Segment.	174
SMCOLUMNS Segment.	176
SMAUDIT Segment.	177
Usage Monitoring RDBMS Table Sizing	178
Renamed or Moved Columns	180
B. Time Fields in Resource Management	183
Time Field Definitions	183
C. Using DDL Statements to Create the Internal Tables	189
Overview of Creating the Resource Analyzer and Resource Governor Internal Tables	189
D. Business Rule Language (BRL)	193
Using Business Rule Language (BRL)	193
BRL Factual Information	194
BRL Keywords	203
Pre-defined Variables	207
Internal Functions	215
Customized Rule Examples	237

Preface

This manual describes how to administer Resource Governor. It provides instructions for configuring the Resource Governor usage monitoring facility and explains how to delete extraneous data from the Resource Governor Usage Monitoring and administrative tables. It presents features such as Rule Parameters, Build Rules, and the Business Rule Builder, and describes the various Resource Governor reports and how to access them. It is intended for anyone who needs to monitor, govern, and report from multiple servers in a distributed data warehouse environment.

For the most current information, please refer to the release notes.

How This Manual Is Organized

This manual includes the following chapters:

Chapter/Appendix	Contents
1 What Is Resource Governor?	Introduces Resource Governor and describes its ability to analyze site activity and contribute to performance analysis. Provides an introduction to Resource Governor's usage monitoring, rule parameters, rule building, and reporting features.
2 Configuring and Administering Resource Management	Describes how to configure and administer Resource Management from the Web Console.
3 Rules Files and Knowledge Bases	Explains how to work with business rules.
4 Resource Governor Report Options	Describes the different categories of Resource Governor reports, the individual reports within these categories, and the process for displaying reports.
5 Getting Help	Describes Resource Governor online help and documentation options.
A Administrative Usage Monitoring Tables Column Descriptions	Summarizes the fields in the Resource Analyzer administrative (system) and Usage Monitoring databases.
B Time Fields in Resource Management	Describes the time fields in the specified synonyms.

Chapter/Appendix		Contents
C	Using DDL Statements to Create the Internal Tables	Explains how to use DDL statements to create the Resource Analyzer and Resource Governor internal tables.
D	Business Rule Language (BRL)	Describes how to customize rules and messages for a particular user or situation using a language called Business Rule Language (BRL).

Conventions

The following table describes the conventions that are used in this manual.

Convention	Description
<p><code>THIS TYPEFACE</code></p> <p>or</p> <p><code>this typeface</code></p>	Denotes syntax that you must enter exactly as shown.
<i>this typeface</i>	Represents a placeholder (or variable) in syntax for a value that you or the system must supply.
<u>underscore</u>	Indicates a default setting.
<i>this typeface</i>	Represents a placeholder (or variable), a cross-reference, or an important term. It may also indicate a button, menu item, or dialog box option that you can click or select.
Key + Key	Indicates keys that you must press simultaneously.
{ }	Indicates two or three choices. Type one of them, not the braces.
[]	Indicates a group of optional parameters. None are required, but you may select one of them. Type only the parameter in the brackets, not the brackets.
	Separates mutually exclusive choices in syntax. Type one of them, not the symbol.

Convention	Description
...	Indicates that you can enter a parameter multiple times. Type only the parameter, not the ellipsis (...).
. . .	Indicates that there are (or could be) intervening or additional commands.

Related Publications

Visit our Technical Content Library at <http://documentation.informationbuilders.com>. You can also contact the Publications Order Department at (800) 969-4636.

Customer Support

Do you have any questions about this product?

Join the Focal Point community. Focal Point is our online developer center and more than a message board. It is an interactive network of more than 3,000 developers from almost every profession and industry, collaborating on solutions and sharing tips and techniques. Access Focal Point at <http://forums.informationbuilders.com/eve/forums>.

You can also access support services electronically, 24 hours a day, with InfoResponse Online. InfoResponse Online is accessible through our website, <http://www.informationbuilders.com>. It connects you to the tracking system and known-problem database at the Information Builders support center. Registered users can open, update, and view the status of cases in the tracking system and read descriptions of reported software issues. New users can register immediately for this service. The technical support section of <http://www.informationbuilders.com> also provides usage techniques, diagnostic tips, and answers to frequently asked questions.

Call Information Builders Customer Support Services (CSS) at (800) 736-6130 or (212) 736-6130. Customer Support Consultants are available Monday through Friday between 8:00 a.m. and 8:00 p.m. EST to address all your questions. Information Builders consultants can also give you general guidance regarding product capabilities. Please be ready to provide your six-digit site code number (xxxx.xx) when you call.

To learn about the full range of available support services, ask your Information Builders representative about InfoResponse Online, or call (800) 969-INFO.

Information You Should Have

To help our consultants answer your questions most effectively, be ready to provide the following information when you call:

- Your six-digit site code (xxxx.xx).
- Your iWay Software configuration:
 - The iWay Software version and release. You can find your server version and release using the Version option in the Web Console.
Note: The MVS and VM servers do not use the Web Console.
 - The communications protocol (for example, TCP/IP or LU6.2), including vendor and release.
- The stored procedure (preferably with line numbers) or SQL statements being used in server access.
- The database server release level.
- The database name and release level.
- The Master File and Access File.
- The exact nature of the problem:
 - Are the results or the format incorrect? Are the text or calculations missing or misplaced?
 - Provide the error message and return code, if applicable.
 - Is this related to any other problem?
- Has the procedure or query ever worked in its present form? Has it been changed recently? How often does the problem occur?
- What release of the operating system are you using? Has it, your security system, communications protocol, or front-end software changed?
- Is this problem reproducible? If so, how?
- Have you tried to reproduce your problem in the simplest form possible? For example, if you are having problems joining two data sources, have you tried executing a query containing just the code to access the data source?

- Do you have a trace file?
- How is the problem affecting your business? Is it halting development or production? Do you just have questions about functionality or documentation?

User Feedback

In an effort to produce effective documentation, the Technical Content Management staff welcomes your opinions regarding this document. You can contact us through our website, <http://documentation.informationbuilders.com/connections.asp>.

Thank you, in advance, for your comments.

iWay Software Training and Professional Services

Interested in training? Our Education Department offers a wide variety of training courses for iWay Software and other Information Builders products.

For information on course descriptions, locations, and dates, or to register for classes, visit our website (<http://education.informationbuilders.com>) or call (800) 969-INFO to speak to an Education Representative.

Interested in technical assistance for your implementation? Our Professional Services department provides expert design, systems architecture, implementation, and project management services for all your business integration projects. For information, visit our website (<http://www.informationbuilders.com/support>).

What Is Resource Governor?

This manual describes the steps for administering Resource Governor. It provides instructions for configuring the usage monitoring facility of Resource Governor and how to delete extraneous data from the Resource Governor Usage Monitoring and administrative tables. The following features are described in detail:

- Rule Parameters
- Build Rules
- Business Rule Builder

This manual also describes the Resource Governor reports and how to access them.

In this chapter:

- [Resource Governor Overview](#)
 - [Resource Governor Features](#)
 - [Resource Governor Operations](#)
 - [How Resource Governor Works](#)
 - [Resource Governor Usage Monitoring](#)
-

Resource Governor Overview

Resource Governor provides Information Systems (IS) organizations with the ability to easily manage the growing volume and unpredictable nature of ad hoc data access. Together with its partner product, Resource Analyzer, Resource Governor is specifically designed to help IS organizations analyze and control end-user data access, as well as provide site-specific control options.

Since the Web is now the medium of choice for disseminating enterprise information among many corporations, understanding user and usage data is critical. The Web environment exposes systems to more access of data by an increasing number of users. This places increased stress on data warehouses, self-service applications, and corporate communications applications. For this reason, it becomes crucial to tune applications for high performance, high availability, and fast response times.

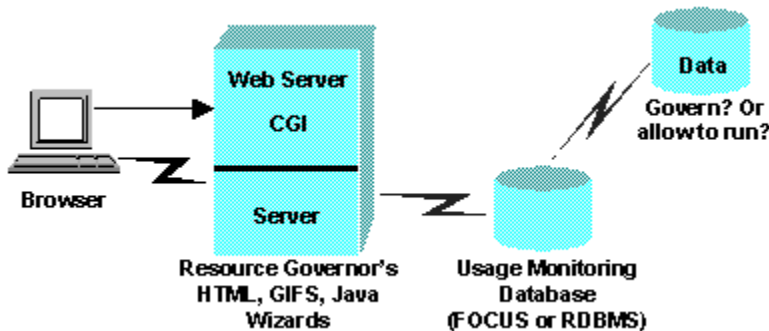
Resource Governor monitors data usage activity by collecting the attributes of requests and storing them in a set of Usage Monitoring tables. An administrator defines site-specific thresholds controlling the amount of resources a request can use. Based on this usage and threshold information, Resource Governor builds rules on how to govern requests against specific data sources. The governing facility uses the rules when inspecting each request, stopping any request that is estimated to exceed the predetermined resource thresholds.

Resource Governor controls monitoring, system configuration parameters, and governing rules. It provides preemptive governing for requests issued to both relational and non-relational data sources.

You can use Resource Governor to monitor, govern, and report from multiple servers in a distributed data warehouse environment.

Note: We use the term *MVS* to refer to z/OS - PDS deployment and *z/OS* to refer to z/OS - HFS deployment using UNIX System Services.

The following diagram illustrates Resource Governor as it resides in the server environment on all platforms:



Resource Governor Features

To help you manage your site activity, Resource Governor controls data access by placing resource limits on requests. To establish this control, Resource Governor:

- Predicts resource usage, relative to the current threshold limits, before executing a request.
- Allows requests within acceptable usage limits to proceed.
- Prevents users from processing requests that are judged to exceed the specified limits.
- Governs data access based on site-specific criteria through business rules and exceptions.

Resource Governor Operations

The following table lists and describes the basic operations of Resource Governor and the steps you need to take to use them. A detailed explanation of each operation is provided later in this manual.

Operation	Description
Collect data about the system.	Turn on usage monitoring for your site. Resource Governor then records the requests and the data about resource usage.
Describe resource thresholds for various time periods.	Enter the relative resource thresholds to set limits on the amount of resources each request can use. Monitored resources include elapsed and CPU time, IOs, and the number of result rows returned to the client. These thresholds help determine which requests can run during a specific time period.
Create rules.	Use Resource Governor to create a rule file. Resource Governor then uses these rules to evaluate each request and predict whether the request operates within the thresholds specified.
Turn Governing on.	Put governing rules into effect for each data source governed.
Govern data source access.	Every time a request uses a monitored data source, Resource Governor predicts whether the request exceeds thresholds based on the rules it has created. The only requests that are allowed to run are those estimated to use lower or equal amounts of threshold time, result rows, and IOs. Requests that are estimated to exceed the thresholds are canceled.
Apply site-specific business rules to add query limitations.	Customize governing to prevent or allow data access under specific conditions.

How Resource Governor Works

Resource Governor monitors request activity, logging the attributes of each request and storing them in a usage-monitoring table. Once an adequate amount of usage monitoring data has been collected, you can run Resource Governor reports, which illustrate various aspects of request activity at your site. The process of choosing resource thresholds and building governing rules can then begin. Resource thresholds must be set before you can build automated or data rules. The rule parameters feature enables you to create shifts, which facilitate the application of date and time restrictions. Once those limits have been set, you can begin to build rules upon which data access will be governed.

At times, you may encounter certain conditions contained in requests that should never be allowed to run. In other cases, you may know that all queries should be allowed to run, such as those submitted by a user with greater privileges. On these occasions, you can create business rules with the Business Rule Builder to check for these kinds of conditions in order to generate cancellation messages, or allow beyond-threshold queries to run. When using business rules only, the selection of resource thresholds is not required. All data sources can be governed by data rules, business rules, or both.

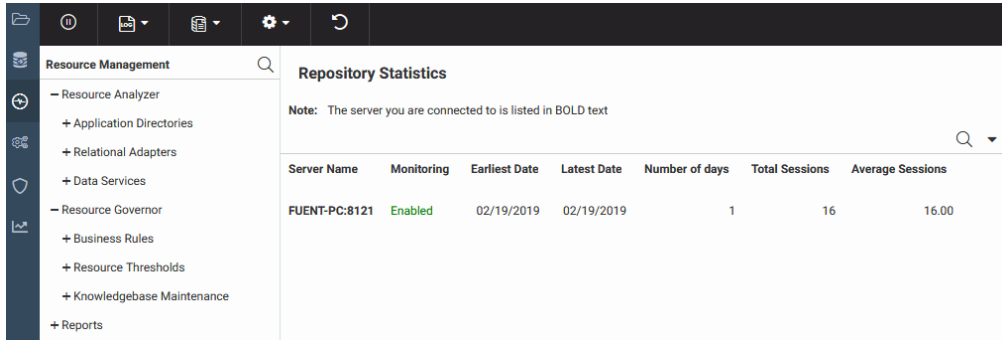
You access the Resource Governor administrative and reporting features through the Web Console of the server.

An alternative form of governing that becomes available as soon as the product is installed, and does not require any monitoring of data requests, is governing with business rules only. Business rules syntax can be created with a New Business Rule while Add Rule for Governing can create a Rule File of Business Rule Type. This rule file may be assigned to any or all data sources, as well as globally for any data request.

Resource Governor Usage Monitoring

Before Resource Governor can monitor request information, it must know your monitoring requirements. You specify these requirements through the Resource Governor Administrator in the Web Console. You can specify a range of monitoring configurations depending on your needs.

The following image shows the Resource Management Console.



The screenshot displays the Resource Management Console interface. On the left is a navigation pane with a search bar and a list of categories: Resource Analyzer, Application Directories, Relational Adapters, Data Services, Resource Governor, Business Rules, Resource Thresholds, Knowledgebase Maintenance, and Reports. The main area is titled 'Repository Statistics' and includes a note: 'Note: The server you are connected to is listed in BOLD text'. Below the note is a table with the following data:

Server Name	Monitoring	Earliest Date	Latest Date	Number of days	Total Sessions	Average Sessions
FUENT-PC:8121	Enabled	02/19/2019	02/19/2019	1	16	16.00

The monitoring process works as follows:

1. A request is issued by a user to access data from a data source.
2. The request is processed by the server, which calls the Usage Monitor prior to retrieving the first record for the request.
3. The Usage Monitor captures the request attributes and stores them until the request finishes retrieving data.
4. The Usage Monitor is called once more to gather resource usage statistics.
5. The logged information is stored in the Usage Monitoring tables. For details about Usage Monitoring fields, see [Administrative Usage Monitoring Tables Column Descriptions](#) on page 157.

For more information on Usage Monitoring, see [Rules Files and Knowledge Bases](#) on page 103.

Due to changing usage patterns, over time you may find that usage-monitoring data that has accumulated is no longer relevant to your current activity and analysis. When this is the case, Resource Governor allows you to delete this data from the usage monitoring table, by using the Resource Governor Administrator.

For more information about using the Resource Governor Administrator to set monitoring configurations and delete data, see [Rules Files and Knowledge Bases](#) on page 103.

Setting Rule Parameters With Resource Governor

Before Resource Governor can create the rules that establish limits on the amount of resources each request can use, you must provide information about permissible resource thresholds at your site and the shifts that control requests made by end users. Resource Governor uses these thresholds and shifts, along with the previously collected usage monitoring data, to build rules about how requests will be governed against specific data sources. To create these thresholds, you use Rule Parameters. Based on these rules, the Resource Governor governing facility decides if a request can execute or not. For more information, see [Rules Files and Knowledge Bases](#) on page 103.

Building Rules

Before Resource Governor can begin governing queries, you must first generate the rules to be used for that process. The rules generated by Resource Governor are based on usage monitoring data and the types of requests that have been monitored before governing is turned on. Business rules may be added without any prior monitoring of requests. When a request is issued, Resource Governor examines the request and uses the rules to determine whether that request should be permitted to execute. For more information about creating, editing, or deleting rules, see [Rules Files and Knowledge Bases](#) on page 103.

Creating Business Rules

You may be aware of certain conditions under which a request should never be allowed to run. However, if these conditions do not occur during the representative sampling of your site applications, they will not become part of the automated Resource Governor rules. Likewise, you may know of certain conditions under which all queries should be allowed to run, such as a user ID with greater privileges. The Business Rule Builder enables Resource Governor to govern immediately on conditions that are known to be unacceptable or that should never be canceled. It also lets you designate more specific parameters beyond simply the time and row number thresholds specified in the Rule Parameters view.

For example, you can configure Resource Governor to govern based on the particular data source being queried, the user ID submitting the request, or even a particular field within the request. Business rules allow for flexibility upon configuration of Resource Governor or changing Information System conditions at your site. Resource Governor uses its proprietary Business Rule Language (BRL) when building rules. For more information about BRL, see [Business Rule Language \(BRL\)](#) on page 193.

For more information on creating business rules, see [Working With Business Rules](#) on page 104.

Reporting With Resource Governor

Resource Governor offers several types of reports, which show, among other things, a summary of monitoring data, a log of canceled queries, and rules.

You can re-sort reports by different columns and drill down on various hyperlinks to view more detailed information and different reports.

You can also specify date-range selection criteria for the reports by accessing a calendar through the Reports Preferences page. These reports are accessed from the Web Console. For more information on each individual report, see [Resource Governor Reports](#) on page 147.

The following is a list of the available reports:

- Monitored Sessions.** Shows the number of connections and amount of resources used.
- Monitored Commands.** Shows the number of commands, as well as how many were used and how many resources are consumed.
- Data Sources Never Used.** Shows the list of application directories with MASTER file names, as well as the last date and time they were used.
- Procedures Never Used.** Shows the list of application directories with FOCEXEC file names, as well as the last date and time they were used.

The Resource Governor Reports folder contains the following reports:

- Monitored Data Sources.** Shows the number of data sources, as well as how many were used and how many resources are consumed.
- Monitored Procedures.** Shows the number of times and what procedures, or FOCEXECs, are used with the resources consumed.
- Resource Thresholds.** Shows parameter details of all Resource Governor thresholds that have been built.
- Knowledgebase Properties.** Shows parameter details of all the knowledgebases that have been built.
- Threshold Advisor.** Recommends resource thresholds based on a statistical model.
- Governed Requests.** Shows all requests canceled by Resource Governor.

For more information on running and viewing these reports, see [Resource Governor Reports](#) on page 147.

Chapter 2

Configuring and Administering Resource Management

The Web Console enables you to configure and administer Resource Management.

In this chapter:

- ❑ [Resource Management Interface](#)
- ❑ [Configuring Resource Management](#)
- ❑ [Administering Resource Management From the Web Console](#)
- ❑ [Migrating the Resource Management Repository](#)
- ❑ [Event Routing](#)


Resource Management Interface

The Resource Management Web Console functionality is available through a ribbon at the top of the interface. The Resource Management ribbon replaces traditional toolbars with a single, rectangular region that spans the width of the application window.

There are two ribbons in Resource Management. The following image shows the ribbon before configuration has taken place.



The following image shows the default ribbon in Resource Management that can be seen once configured.



Ribbon

The Resource Management ribbon represents a subset of controls and commands that are organized into related groups depending on function or relationship. All groups are labeled and all icons represent a control or command.

From the ribbon you can perform actions such as enabling or disabling monitoring, archiving logs, and setting logging properties. For more information on all of the commands shown in the ribbon, see [Resource Management Configuration Maintenance](#) on page 35.

Configuring Resource Management

At run time, Resource Management collects usage monitoring data in a log file. At predetermined times, such as during server startup or a user request, the log data collected is archived to a permanent repository. Before collecting the data, Resource Management must be configured.

To configure Resource Management, after the server is installed, you must:

1. Consider security authorization for RM files and libraries.

There are two levels of security to consider:

- Administrator:** MVS requires read and write authority to all datasets allocated to EDACCFG, EDACMAS, EDACACX, and SMARTLIB DD statements of the server. All other platforms require read and write authority to EDACONF/catalog and EDACONF/catalog/rm.
- Basic user:** MVS requires read authority to the dataset allocated to the SMARTLIB DD statement.

Note: In order for User activity to be collected correctly by Resource Management, all users who connect to the WebFOCUS Reporting Server need to have Executive privilege for the EDACONF/catalog/rm directory. The default directory name is ibi\srvnn\wfs\catalog\rm (for Windows) or ibi/srvnn/wfs/catalog/rm (for UNIX), where *nn* is the Reporting Server release.

While the standard security roles, Server Admin, Application Admin, Operator, and Basic User, have this privilege available to them after a server install, any Custom roles need to be explicitly configured for EDACONF/catalog/rm with the Execute (PRRUN) privilege.

2. Enable Resource Management from the Web Console.
3. Type your license key and select a repository type. Resource Management can be configured to create Usage Monitoring tables in a FOCUS repository or in a relational repository. The repository holds data on resource usage.
4. Configure the repository.

Procedure: How to Configure Resource Management Using a FOCUS or Relational Repository

To enable Resource Management, you must have a valid license key. After typing the key, you can specify the type of repository to be used for Usage Monitoring. The Usage Monitoring tables are then created. This procedure shows how to create Usage Monitoring tables in a FOCUS or relational repository using a license code for Resource Governor.

Before performing the following steps, make sure the server is running.

1. On a Windows system, click the *Start* menu, and then click *Programs*. If using Windows 10, click *All Apps*.
2. Select *iWay Software* (or *Information Builders*, if you have a WebFOCUS server installation), then the server type, and click *Web Console*.

If you are using an operating system other than Windows, enter

`http://ipaddress:port`

where:

`ipaddress`

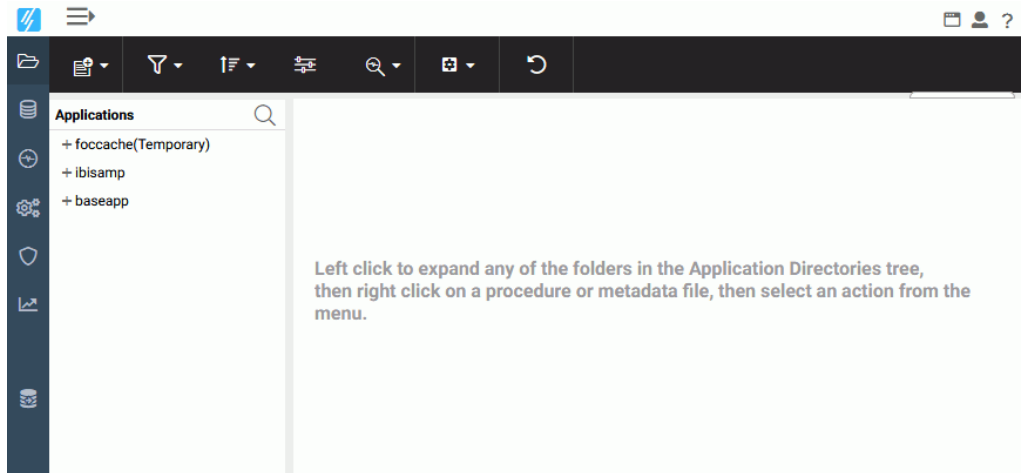
Identifies the server where Resource Governor is to be configured. Do not use *localhost* when accessing the Web Console for Resource Management initial configuration.


`port`

Is the port number for that server.

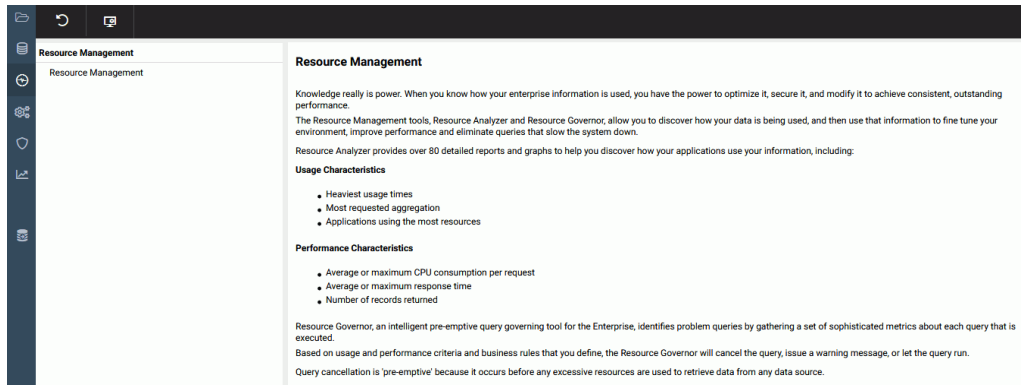
Note: On the Programs or All Apps menu under iWay Software, the server types can include iWay Full-Function Server, or iWay DataMigrator Server. On the Programs or All Apps menu under Information Builders, the server type can include a WebFOCUS Server. Resource Governor can be configured on any of these servers.

The Web Console opens, as shown in the following image.



3. Click the *Resource Management* button  on the sidebar.

The Resource Management page opens, as shown in the following image.



4. Click *Configure* on the ribbon.
5. Type the license key from the packing slip. There are different keys: one for Resource Analyzer, one for Resource Governor, and one if both products are licensed.
6. Click *Next*.

The Resource Management Configuration pane opens.

7. For a FOCUS Repository:

- a. Type a Server Name, which is used by Resource Governor as the server name in collected data and on reports. Any identifier of up to 48 characters is acceptable, but the value should be unique across all servers where Resource Governor is installed. The default value is taken from the browser address bar, and is usually the computer name, along with the port number used by the server. It is recommended that this value is not changed. If *localhost:port number* is used in the browser, it can not be used for the Server Name and must be changed.

Note: If you need to change the server name after configuring Resource Management, you have to remove the current configuration and reconfigure it using the new server name. Note that removing the current configuration of Resource Management will delete all of your collected data.

- b. Select *FOCUS* from the Repository drop-down menu.
- c. Go to step 8g.

8. For a Relational Repository:

- a. Type a Server Name. Used by Resource Governor as the server name in collected data and on reports. Any identifier up to 48 characters is acceptable, but the value should be unique across all servers where Resource Governor is installed. The default value is taken from the browser address bar, and is usually the computer name, along with the port number used by the server. It is recommended that this value is not changed. If *localhost:port number* is used in the browser, it can not be used for the Server Name and must be changed.

Note: If you need to change the server name after configuring Resource Management, you have to remove the current configuration and reconfigure it using the new server name. Note that removing the current configuration of Resource Management will delete all of your collected data.

- b. Select *Relational* from the Repository drop-down menu. The following image shows the additional fields that become available.

Configure Resource Management

Repository Settings
? Server Name
? Repository
? Data Adapter
? Create Repository
? Owner/Schema (case sensitive)

Configuration Settings
? Data Request Monitoring
? Procedure Request Monitoring
? WebFOCUS System Procedure Monitoring
? Repository load type (Before selecting Bulk load confirm that the vendor's utility has been installed)

Note: If no certified Resource Management adapter has been configured, the only option in the drop-down menu will be FOCUS.

- c. Select an adapter from the Data Adapter drop-down menu. This adapter identifies the data adapter and the connection name to be used for storing Resource Governor Usage Monitoring data. The connection name is the logical name used to identify a particular set of connection attributes. This setting also determines the SUFFIX used when creating the Master Files for these tables.

The current list of certified relational engines is:

- Db2
- Informix
- MariaDB
- MS SQL Server
- MySQL
- Oracle
- PostgreSQL
- Sybase
- Teradata
- Vertica

- Generic ODBC/JDBC for Sybase SQL Anywhere Release 12 non-Unicode only.

Note:

- The Sybase ASA database used for the Resource Management repository must be created with the *Ignore trailing blanks in string comparisons* option (or the `-b` option in the `dbinit` command).
- If Db2 is going to be the target for the Resource Management repository, the DBA of the site should investigate the use of the parameter `LOCKSIZE` and change it to `ANY` if large volumes of data are to be stored. This value of `LOCKSIZE` will prevent the Db2 error message:

The limit on the number of locks has been exceeded for the tablespace within Db2.

If the DDL Only option is used at configuration time, then this parameter can be added to the DDL that is generated. Otherwise, the DBA of the site can add it to the existing Resource Management environment.

When Resource Management data is archived to a target DB2 repository, the user ID that is used to submit the archive job must have the authority to create temporary tables in DB2 with a minimum page size of 4K.

- d. Set Create Repository to Yes, unless the repository has already been created using the native Data Definition Language (DDL) of the data adapter or you are configuring a many-to-one environment.

When No is selected, the configuration process will not issue `DROP TABLE`, `DROP INDEX`, `CREATE FILE`, and `CREATE INDEX` commands. Because the configuration process will always initialize the repository with default data and issue the default `GRANTS`, the repository must exist when the commands are executed or the configuration will fail. Yes is the default value.

For more information, see [Using DDL Statements to Create the Internal Tables](#) on page 189. For more information on many-to-one, see [Configuring Many-to-One Monitoring](#) on page 30.

- e. Provide an *Owner/Schema*. By default, the tables are created for the specified owner/schema in the default database, and the owner/schema must be authorized to create tables in the database.

If the tables are to be created in an alternate database that the owner has rights to, and the technique is supported by the RDBMS, the database name must be specified as well (for example, `database.owner`). Consult the appropriate RDBMS documentation for the appropriate syntax. On the System i platform, the Owner Name will be shown as Library.

- f. For servers running on z/OS, HFS, or PDS deployment where Db2 has been selected as the data adapter, provide a *Db2 DBSPACE*. The format of the input required is *Database* or *Database.tablespace*.
- g. Set Global Settings, which has the following options:

Data Request Monitoring

Indicates whether Data Request Monitoring is on or off. When Global Data Request Monitoring is enabled, or set to *On*, Resource Governor collects information on all SQL SELECT and FOCUS TABLE, TABLEF, MATCH, GRAPH, MODIFY, and MAINTAIN requests. The default value is *Off*. Monitoring will be by individually selected data sources.

Note: On initial configuration it is recommended that this value is set to *On* so that all default monitoring data can be collected for an initial period after which reports can be run to determine at what level future monitoring should be set. Default monitoring is down to the Field/Column level, which could amount to a large volume of Repository data. If, after reviewing the available reports, it is determined that, for example, this level of monitoring is not required, it can be turned off. For more information, see [How to Set Repository Properties](#) on page 63.

Procedure Request Monitoring

Indicates the level of Procedure Request Monitoring. The following options are available:

- Execute and Include.** All executed procedures and any procedure that used -INCLUDE will be monitored. This is the default value.
- Execute Only.** Any procedure that is executed using EX or EXEC will be monitored.
- Primary Execute Only.** The only procedure monitored will be the primary procedure. Any procedure executed from within the primary (-INCLUDE or EX) will not be monitored.

WebFOCUS System Procedure Monitoring

Determines whether WebFOCUS internal system procedures will be monitored. The following options are available:

- Off.** Internal system procedures will not be monitored. This is the default setting.
- On.** Internal system procedures will be monitored.

Repository load type

Indicates the load method that will be used to add log data to the data repository. The following options are available:

- Bulk load.** The bulk load utility for the selected RDBMS server, if available, will be used to load the data into a temporary table on the RDBMS server using the connection adapter that was selected for the repository. This is the default value.
- Table load.** The TABLE command will be used to load the data into a temporary table on the RDBMS server using the connection adapter that was selected for the repository.

After the data is loaded into the temporary table using either load method, the MERGE command is then used to merge the data from the temporary table into the data repository.

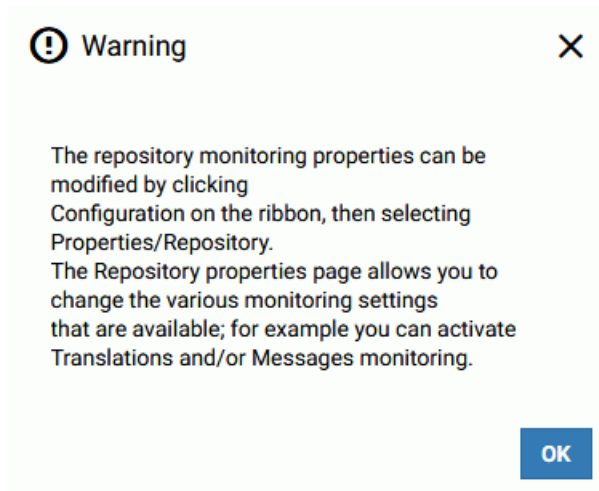
9. Click *Configure*.

The Repository Statistics page opens, and the navigation pane appears on the left, as shown in the following image.

Server Name	Monitoring	Earliest Date	Latest Date	Number of days	Total Sessions	Average Sessions
FUENT-PC:8121	Enabled	03/05/2018	03/06/2018	2	213	106.50

Note: The server that you are connected to is shown in bold text.

After a successful configuration, a confirmation dialog box opens, as shown in the following image.



You can now select monitoring levels for Resource Governor. For more information, see [How to Set Repository Properties](#) on page 63.

Click *OK*.

Configuring Many-to-One Monitoring

Several considerations and steps must be taken in order to configure Resource Management to gather monitor data for multiple servers into a single repository. This can only be done using relational adapters. Access must be available from each server machine to the selected relational engine, and a license will be needed for each individual adapter. All servers must be at the same major release level.

All servers in a cluster environment, or multiple servers on UNIX, IBM i/MVS/PDS, and Windows platforms, can collect monitor data into a single repository. If you have environments with multiple WebFOCUS or iWay servers configured for Resource Management and want to collect monitor data from those servers into a single repository, the steps in the following procedure show how to configure Resource Management using the Web Console of the servers.

Configuring Resource Governor

If you have not used the Resource Management DDL to create the repository tables, then you need to pick a server to do the initial configuration that will create the repository tables. Using the Web Console of the server, you will need to perform the steps in [How to Configure the First Server](#) on page 31. If the Resource Management DDL was used to create the repository tables, then go to [How to Configure Additional Servers](#) on page 31.

Procedure: How to Configure the First Server

For the first server, configure Resource Management from the Web Console of the server using the following steps:

1. Configure the relational data adapter for the repository, if it does not yet exist.
2. Configure Resource Governor from the Resource Management Configuration pane.
3. Enter the Resource Management license key. Click *Continue*.
4. Select *Relational* from the Repository Type drop-down menu.
5. Enter a unique name in the Server Name field, or accept the default.
6. Select a Data Adapter and the appropriate connection name from the drop-down menu.
7. Set Create Repository to Yes.
8. Provide an Owner/Schema name for the relational adapter that Resource Governor will use to create the repository tables.
9. Click *Configure*.

After configuration is complete, select the monitoring level for Resource Governor. For more information, see [How to Set Repository Properties](#) on page 63.

Procedure: How to Configure Additional Servers

For additional servers that will use the same repository, configure Resource Management from the Web Console of the server using the following steps:

1. Configure the relational data adapter for the repository, if it does not yet exist, that connects to the same database used in step 1 in [How to Configure the First Server](#) on page 31.
2. Configure Resource Governor from the Resource Management Configuration pane.
3. Enter the Resource Management license key. Click *Continue*.
4. Select *Relational* from the Repository Type drop-down menu.
5. Enter a unique name in the Server Name field, or accept the default.
6. Select a Data Adapter and the appropriate connection name from the drop-down menu.

7. Set Create Repository to *No*.

Note: Create Repository must be set to *No* on each additional server. If set to *Yes*, an error message will be displayed stating that the repository already exists.

If the relational repository tables were created by the DBA using the DDLs provided with the product, then all servers have to be configured using Create Repository *No*.

8. Provide the same Owner/Schema name for the relational adapter that Resource Governor used to create the repository tables.
9. Click *Configure*.

After configuration is complete, select the monitoring level for Resource Governor. For more information, see [How to Set Repository Properties](#) on page 63.

Note: All administration for Resource Governor must be done on each of the servers configured to use the same repository tables.

Reference: Governing

If Governing is being used, administration for Resource Governor must be done on each server. Knowledgebase rule files cannot be shared or copied from one server to another.

When rules are built for Governing, there is a knowledgebase rule file built on the server. These knowledgebase files can only be used on the server that they were built on. For example, if rules to govern databases are built on a UNIX server, these rules can only be used on that specific UNIX server. The same rule applies for Windows, z/OS, PDS server, and other platforms.

Administering Resource Management From the Web Console

The Resource Governor Usage Monitoring facility allows you to monitor requests as they are submitted by users, and logs information about request usage. This information includes the usage statistics associated only with data retrieval, or data manipulation in the case of MODIFY or MAINTAIN. This request information is stored in Resource Governor Usage Monitoring tables. You can specify the Monitor Preference, which allows you to change what information is collected for monitored data sources.

The Web Console enables you to administer Resource Management.

You can:

- Enable/disable monitoring and governing.
- Maintain the repository.

- Reconfigure the repository from FOCUS to relational, or one relational environment to another relational environment.
- Remove Resource Management.
- Enable/disable global monitoring.
- Set monitoring preferences.
- Add data sources to be monitored.
- View the properties of data sources and the sample data they produce.
- Monitor Data Services.
- Access reports.


Access to these tasks is provided by the Resource Management button on the sidebar.

The following image shows the hierarchical navigation menu that is available when the Resource Management button is selected.

- Resource Analyzer
 - + Application Directories
 - + Relational Adapters
 - + Data Services
- Resource Governor
 - + Business Rules
 - + Resource Thresholds
 - + Knowledgebase Maintenance
- Reports
 - 🔍 Monitored Sessions
 - 🔍 Monitored Commands
 - 🔍 Repository Statistics
 - 🔍 Data Sources Never Used
 - 🔍 Procedures Never Used
 - 🔍 ACI Monitoring
- Resource Governor
 - 🔍 Monitored Data Sources
 - 🔍 Monitored Procedures
 - 🔍 Resource Thresholds
 - 🔍 Knowledgebase Properties
 - 🔍 Threshold Advisor
 - 🔍 Governed Requests

Resource Management Configuration Maintenance

The Resource Management ribbon provides options for configuring Resource Management, as shown in the following image.



The following sections provide detailed explanations for each option in the ribbon.

Enabling or Disabling Monitoring

When you enable Resource Management, it automatically turns on monitoring. However, you can disable and re-enable monitoring at any time.

Procedure: How to Enable or Disable Monitoring

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Enable/Disable* icon.

When disabling Resource Management, a confirmation window will be displayed. Click *OK* to disable Resource Management or *Cancel* to leave it enabled.

Note: In prior releases of the WebFOCUS Reporting Server, the following line was added to the EDASPROF.prf file when Resource Management is configured and enabled.

```
SET SMARTMODE=ON
```

It was possible to remove this line manually to disable the product rather than use the option detailed above. That SET command is no longer used and the only way to Disable/Enable the product is as documented above.

Logging

The Reporting Server will write all information regarding monitored sources to a log file. This log file is then archived to the Resource Management repository at predetermined intervals or limits, while the server is running. The current log is also archived when the server starts or when the administrator submits a job from the Web Console of a running server.

The intervals or limits can be configured from the Logging option in the Resource Management Web Console. When configured, the log file of an active server will be archived based on the maximum number of sessions or the days set, or based on a scheduled job configured by the administrator.

Procedure: How to Set Logging Properties

Logging properties can be configured in the Resource Management Web Console. When configured, the log files will be archived based on the maximum number of sessions and days set, or based on a scheduled job configured by the administrator.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Logging*, and then click *Properties*.

The Resource Management Logging Properties pane opens, as shown in the following image.

Logging Properties

Repository archiving based on number of sessions and days will not be supported in the next release. Please switch to Schedule Archive prior to the next release.

Maximum sessions per log	<input type="text" value="10000"/>	Values: 500 - 999999
Maximum days per log	<input type="text" value="1"/>	Values: 1 - 365
Schedule Archive	<input type="button" value="No"/>	
Limit Agent memory	<input type="button" value="No"/>	

The following options are available:

Maximum sessions per log

The number of sessions recorded in the active log file before the log is archived and a new log is started. The default value is 10000. This option is only available if Schedule Archive is set to *No*.

Maximum days per log

The number of days an active log file is kept open before it is archived and a new log is started. The default value is 1. This option is only available if Schedule Archive is set to *No*.

Schedule Archive

When set to *Yes*, only the Limit Agent memory option for interim log records is shown and the administrator has to configure the schedule archive settings. For more information, see [How to Schedule an Archive](#) on page 39.

When set to *No*, automatic archiving is active and based on the two additional settings. The default value is *No*.

Note: To use the scheduler feature for archiving the log file, the server Scheduler process must be running. To check this, from the Web Console, click the *Workspace* option on the sidebar. Expand the *Special Services and Listeners* folder, right-click *SCHEDULER*, and click *Start*. If the *Start* option is not available, the scheduler is already running. If Resource Management and setting Schedule Archive to *Yes* are done in the same Web Console session, the Scheduler will need to be restarted. To do so, right-click *SCHEDULER* and click *Stop*, followed by *Start*.

Limit Agent memory

This setting is used for a connection during which multiple procedures are run within the session. This type of connection can cause the Resource Management logging process to use a large amount of memory before the log records are written to the physical disk file (*rmldata.log*).

When set to *Yes*, enter a value between 5 and 99 (in megabytes). When this parameter is set to *Yes* and the specified amount of memory is reached, the connected user log records in memory are written to a temp disk file and the memory is freed.

At this point any new log records generated will be written to disk instead of memory until the end of the session. At this time, all the log records in the temporary disk file are written to the physical disk file (*rmldata.log*). This setting should only be used if excessive memory usage is observed and users are connected for long periods.

4. Make your desired changes and click *Save*.

Procedure: How to View Logs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Logging*, and then click *View Logs*.

The Data log list pane opens, as shown in the following image.



File Name	Size (KB)	Last Modified
rmldata.log	38	2018/03/05 15.58.02

If you right-click any item in any column, the following options are available:

- View.** Opens the log file for the selected data log. You can also double-click the log name to view the file.
- Download.** Saves the selected log to your Downloads folder. In Windows, the default location of this folder is C:\users*username*\Downloads.

Archive Log Files

When this option is selected, a context menu is shown with the following options:

- Schedule and E-Mail

Note: This option is only available when Schedule Archive is set to Yes. For more information on logging properties, see [How to Set Logging Properties](#) on page 36.

- Submit job
- View jobs

These give the administrator the option to archive the current log file (a new log file is automatically started), as well as view previous archive jobs, including any that were submitted due to server restart or any of the intervals or limits mentioned above having been reached. For more information, see [How to Submit an Archive of the Current Log File](#) on page 46.

Note: After a successful archive, the log file is deleted. If the archive fails, the log file will remain available for archiving. The failure should be investigated and corrected. Once this is done, the *List Unarchived Log Files* option can be used to archive the data. For more information, see [List Unarchived Log Files](#) on page 47.

Reference: Bulk Load Support

The archive process uses two steps. The first step will use the Bulk Load facility to move the monitored data to a staging/temporary table and then issue a SQL MERGE syntax to insert the data into the repository. This means that the adapter connection to the repository environment must have correct privileges to create a table.

If the RDBMS repository vendor supported by Resource Management supports bulk load and the environment is correct, the archive process will attempt to issue bulk load commands. If they fail, the archive will fall back to previous methods of inserting monitored data.

Bulk load is transparent except for the following:

- Microsoft SQL Server on Windows.** This requires the BCP utility to be available in the windows search PATH. Bulk load is not available on non-Windows platforms.

Procedure: How to Schedule an Archive

Note: Only one archive can be scheduled at a time.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive, Schedule and E-Mail*, and click *Manage jobs*.

The Manage Schedule and E-Mail for procedure pane opens, as shown in the following image.

^ General
? Schedule Status No Schedule ▾ Enables/Disables Schedule

▽ Log and output destinations

▽ E-Mail On Start

^ E-Mail On Completion

? Enable

? Mail to To

Cc

Bcc

Use a semicolon as the separator character for multiple email addresses

? Advanced

▽ E-Mail On Failure

Cancel Set

The Manage Schedule and E-Mail for procedure pane contains the following fields and options:

Schedule Status

Determines if there is a schedule and if so, whether the schedule is active or inactive.

Log and output destinations

These options control where the log information from the Maintenance Job should be sent when the Job finishes.

The following options are available:

- ETLLOG.** The log is sent to the DataMigrator log.
- EDAPRINT.** The log is sent to the Server log.
- E-Mail.** The log is sent as an attachment to an email message if specified in Job Properties.

- Send HTML output as inline message.** When selected, if any stored procedures called by the flow generate an output, the message is included as message text instead of as an attachment. This option is only available in the E-Mail On Completion section.

Multiple options can be selected. ETLLOG and E-Mail are selected by default.

Note:

- If the job is run from the DMC or Web Console, information will go to the console log regardless of this setting.
- The *E-Mail On Start* and *E-Mail On Completion* options are not relevant to the Resource Management Archive job and should not be used.

Enable

Select this to enable when an email should be sent.

Mail to

The email addresses where the messages will be sent. Place each email address (or user ID configured on the server with an associated email address) on a new line with no separators.

Advanced

Select this option to get the following options:

- Mail to addresses defined in procedure.** You can send the messages to email addresses defined in the procedure by either typing in the location of the procedure, or clicking the ellipsis button and manually selecting it from the Select Procedure dialog box.
- Importance.** Select the importance for the e-mail from the drop-down menu. Choices are Low, Normal, or High.
- Subject.** The subject line of the message.
- E-Mail Message.** The body of the message.

- If you select *Active* or *Inactive* from the Schedule Status drop-down menu, the Manage Schedule and E-Mail for procedure pane opens with more fields, as shown in the following image.

^ General	
? Schedule Status	Active <input type="button" value="v"/> Enables/Disables Schedule
? Schedule Type	Run Once <input type="button" value="v"/> Determines whether a flow will execute
^ Start and stop range	
? Start Date	05/07/2018 <input type="button" value="calendar"/> Specifies the start time and date for the execution of
? Start Time	12 <input type="button" value="v"/> 27 <input type="button" value="v"/> Specifies the start time and date for the execution of
^ Special date/time ranges	
? Specify Stop Time	<input type="checkbox"/> Specifies forced end time
v Log and output destinations	
v E-Mail On Start	
^ E-Mail On Completion	
? Enable	<input type="checkbox"/>
? Mail to	To <input type="text"/>
	Cc <input type="text"/>
	Bcc <input type="text"/>
	Use a semicolon as the separator character for multiple email addresses (ic
? Advanced	<input type="checkbox"/>
v E-Mail On Failure	
<input type="button" value="Cancel"/>	<input type="button" value="Set"/>

The Manage Schedule and E-Mail for procedure pane contains the following additional fields and options:

Schedule Type

Determines whether a process flow will execute once, on a recurring basis, on several specific days (multi-day), or when the server starts.

Interval Type

Specifies the interval in which a process flow will be executed. For example, if you select Month as your Interval Type, and select 2 as your Interval Number, the flow will execute every two months. This option is only available if the schedule type is set to *Recurring*.

Interval Value

Specifies the frequency at which a process flow will be executed. For example, if you select Week as your Interval Type, and select 2 as your Interval Value, the flow will execute every two weeks. This option is only available if the schedule type is set to *Recurring*.

Start Date

Specifies the start date for the execution of the process flow.

Start Time

Specifies the start time for the execution of the process flow using 24-hour time.

Specify Stop Time

Allows you to specify an end time for the execution of the process flow using 24-hour time. This option is only available if the schedule type is set to *Run Once* or *Multi-Day*.

Specify End Date

Allows you to specify an end date and time for the execution of the process flow. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

End Date

Specifies the end date for the execution of the process flow. This option is available if Specify End Date is selected.

End Time

Specifies the end time for the execution of the process flow using 24-hour time. This option is available if Specify End Date is selected.

Days of the Week

Specifies the day of the week on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Days of the Month

Specifies the day of the month on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Intraday Start

Specifies the beginning of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Intraday End

Specifies the end of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Run on Holidays

Specifies the option to execute the process flow on desired holidays.

The following options are available:

- Skip.** Excludes run on desired holidays.
- Only.** Only run on desired holidays.

This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

5. Click **Set**.

A confirmation dialog box opens.

6. Click **OK**.

Procedure: How to View Scheduler Agents

This report runs for the user whose ID value is stored in the `sched_scan_id` of a Scheduler. The Scheduler must be active to run the Scheduler Agents report.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive*, then point to *Schedule and E-Mail*, and click *Scheduler Agents*.

The Scheduler Agents pane opens, as shown in the following image.

ID	State	Application	Procedure	Report	User ID	Schedule Type	Last Run Time	Next Run Time	Tscomid
D7	SCHEDULED	_edaconf/catalog	F rmusavlg		OPSYS\BIVcf12026	Recurring	Not Available	2018/09/06 15:30	--

The following buttons are available on the Scheduler Agents pane:

Choose States

Allows you to select the states that are displayed on the Scheduler Agents pane.

Activate Selected

Allows you to enable the Scheduler for the selected procedure.

Deactivate Selected

Allows you to disable the Scheduler for the selected procedure.

If you right-click a drill-down item in the Procedure column, the following options are available:

Manage Schedule/E-Mail

Opens the Manage Schedule and E-Mail for procedure pane, which allows you to edit the schedule settings for the selected procedure.

View Scheduled Events

Opens the Scheduled Events pane, which allows you to view jobs scheduled to start during a specified time range.

View Log

Opens the Log report for the selected procedure.

View Statistics

Opens the Statistics report for the selected procedure.

Quiesce Scheduler Agent

Allows you to stop the selected scheduled procedure from running.

Procedure: How to View Scheduled Events

Note: The Scheduler must be active to run the Scheduler Events report.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive*, then point to *Schedule and E-Mail*, and click *Scheduled Events*.

The Scheduled Events pane opens, as shown in the following image.

View jobs scheduled to start during specified time range

2018/9/6 2018/9/6

00 ▾ 00 ▾ 23 ▾ 59 ▾

View Report Cancel

4. Set the desired time range and click *View Report*.

The report opens, as shown in the following image.

Application	Procedure	Run Date	Run Time	User ID
_edaconf/catalog	rmusavlg	2018/09/06	15:45:00	OPSYS\IBI\cf12026

This report shows all flows and procedures scheduled to run during the time period you selected. The Run Date and Run Time columns show their scheduled run times. The User ID column shows the user ID used when running the flow or procedure.

Procedure: How to Submit an Archive of the Current Log File

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive*, and then click *Submit job*.

The Warning dialog box opens.

4. Click *OK* to save the log data.

Procedure: How to List Previous Archive Jobs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive*, and then click *View jobs*.

The Deferred List opens showing the saved logs, as shown in the following image.

<input type="checkbox"/> Delete Selected <input type="checkbox"/> Refresh Selected		Defer ID	User ID	Status	App	Procedure	Report	Submitted Time	Scheduled Time	Days Left
<input checked="" type="checkbox"/>		2018-09-05-10-29-06cmrpip000001_LOOPBACK	OPSYS\IBI\cf12026	Ready		RMARCIVE		2018 September 05 16:41:38	---	29
<input type="checkbox"/>		2018-09-05-10-29-06cmrpip000005_LOOPBACK	OPSYS\IBI\cf12026	Ready		RMARCIVE		2018 September 06 12:01:41	---	29
<input type="checkbox"/>		2018-09-05-10-29-06cmrpip000006_LOOPBACK	OPSYS\IBI\cf12026	Ready		RMARCIVE		2018 September 06 15:30:06	---	29
<input type="checkbox"/>		2018-09-05-10-29-06cmrpip000007_LOOPBACK	OPSYS\IBI\cf12026	Ready		RMARCIVE		2018 September 06 15:45:05	---	29
<input type="checkbox"/>		2018-09-05-10-29-06cmrpip000008_LOOPBACK	OPSYS\IBI\cf12026	Ready		RMARCIVE		2018 September 06 15:57:20	---	29

If you right-click an item in the Defer ID column, the following options are available:

- View.** Opens the output of the selected item.
- Download.** Saves the selected log to your Downloads folder. In Windows, the default location of this folder is C:\users\username\Downloads.

- ❑ **Delete.** Deletes the selected item.

List Unarchived Log Files

Select this option to show a list, if any exist, of previous log files that have not been archived.

Procedure: How to List Unarchived Logs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Logs* icon, point to *Archive*, and then click *List Unarchived*.

The Resource Management unarchived logs list opens, listing the non-archived logs.

4. Select an unarchived log and click *Next* to archive the desired unarchived logs.
5. Select a date and time to archive and click *Continue*.

Working With Repository Maintenance

Repository data can be deleted once it has been used or reported on. This deletion can be by date range or, if sufficient data has been stored, the data can be deleted if it is more than one, three, or six months old.

Procedure: How to Schedule a Repository Maintenance

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Maintenance*, then point to *Schedule and E-Mail*, and click *Add Job*.

The Resource Management Schedule Repository Maintenance pane opens, as shown in the following image.

Schedule Repository Maintenance for FUENT-PC:8121

WARNING: The repository tables for this server should be backed up before proceeding.

Delete data older than:

- 1 month
- 2 months
- 3 months
- 6 months
- 12 months
- Specify months

4. If you have multiple servers available, select the server from the drop-down menu.
5. Select the time range of data that you want to schedule.
6. Click *Next*.

The Manage Schedule and E-Mail for procedure pane opens, as shown in the following image.

^ **General**

? Schedule Status Enables/Disables Schedule

^ **E-Mail On Completion**

? Enable

? Mail to **To**

Cc

Bcc

Use a semicolon as the separator character for multiple email addresses

? Advanced

^ **E-Mail On Failure**

The Manage Schedule and E-Mail for procedure pane contains the following fields and options:

Schedule Status

Determines if there is a schedule and if so, whether the schedule is active or inactive.

Log and output Destinations

These options control where the log information from the Maintenance Job should be sent when the Job finishes.

The following options are available:

- ETLLOG.** The log is sent to the DataMigrator log.
- EDAPRINT.** The log is sent to the Server log.
- E-Mail.** The log is sent as an attachment to an email message if specified in Job Properties.
- Send HTML output as inline message.** When selected, if any stored procedures called by the flow generate an output, the message is included as message text instead of as an attachment. This option is only available in the E-Mail On Completion section.

Multiple options can be selected. ETLLOG and E-Mail are selected by default.

Note:

- If the job is run from the DMC or Web Console, information will go to the console log regardless of this setting.
- The *E-Mail On Start* and *E-Mail On Completion* options are not relevant to the Resource Management Archive job and should not be used.

Enable

Select this to enable when an email should be sent.

Mail to

The email addresses where the messages will be sent. Place each email address (or user ID configured on the server with an associated email address) on a new line with no separators.

Advanced

Select this option to get the following options:

- Mail to addresses defined in procedure.** You can send the messages to email addresses defined in the procedure by either typing in the location of the procedure, or clicking the ellipsis button and manually selecting it from the Select Procedure dialog box.
- Importance.** Select the importance for the e-mail from the drop-down menu. Choices are Low, Normal, or High.
- Subject.** The subject line of the message.
- E-Mail Message.** The body of the message.

7. If you select *Active* or *Inactive* from the Schedule Status drop-down menu, the Manage Schedule and E-Mail for procedure pane opens with more fields, as shown in the following image.

^ General	
? Schedule Status	Active <input type="button" value="v"/> Enables/Disables Schedule
? Schedule Type	Run Once <input type="button" value="v"/> Determines whether a flow will execute
^ Start and stop range	
? Start Date	05/07/2018 <input type="button" value="calendar"/> Specifies the start time and date for the execution of
? Start Time	12 <input type="button" value="v"/> 27 <input type="button" value="v"/> Specifies the start time and date for the execution of
^ Special date/time ranges	
? Specify Stop Time	<input type="checkbox"/> Specifies forced end time
^ Log and output destinations	
^ E-Mail On Start	
^ E-Mail On Completion	
? Enable	<input type="checkbox"/>
? Mail to	To <input type="text"/>
	Cc <input type="text"/>
	Bcc <input type="text"/>
	Use a semicolon as the separator character for multiple email addresses (ic
? Advanced	<input type="checkbox"/>
^ E-Mail On Failure	
<input type="button" value="Cancel"/>	<input type="button" value="Set"/>

The Manage Schedule and E-Mail for procedure pane contains the following additional fields and options:

Schedule Type

Determines whether a process flow will execute once, on a recurring basis, on several specific days (multi-day), or when the server starts.

Interval Type

Specifies the interval in which a process flow will be executed. For example, if you select Month as your Interval Type, and select 2 as your Interval Number, the flow will execute every two months. This option is only available if the schedule type is set to *Recurring*.

Interval Value

Specifies the frequency at which a process flow will be executed. For example, if you select Week as your Interval Type, and select 2 as your Interval Value, the flow will execute every two weeks. This option is only available if the schedule type is set to *Recurring*.

Start Date

Specifies the start date for the execution of the process flow.

Start Time

Specifies the start time for the execution of the process flow using 24-hour time.

Specify Stop Time

Allows you to specify an end time for the execution of the process flow using 24-hour time. This option is only available if the schedule type is set to *Run Once* or *Multi-Day*.

Specify End Date

Allows you to specify an end date and time for the execution of the process flow. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

End Date

Specifies the end date for the execution of the process flow. This option is available if Specify End Date is selected.

End Time

Specifies the end time for the execution of the process flow using 24-hour time. This option is available if Specify End Date is selected.

Days of the Week

Specifies the day of the week on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Days of the Month

Specifies the day of the month on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Intraday Start

Specifies the beginning of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Intraday End

Specifies the end of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Run on Holidays

Specifies the option to execute the process flow on desired holidays.

The following options are available:

- Skip.** Excludes run on desired holidays.
- Only.** Only run on desired holidays.

This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

8. Click *Set*.

A confirmation dialog box opens.

9. Click *OK*.

Procedure: How to View Scheduled Jobs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Maintenance*, then point to *Schedule and E-Mail*, and click *List jobs*.

The Schedule Repository Maintenance - Update pane opens, as shown in the following image.

Schedule Repository Maintenance - Update

Name	Status	Months	Description
rmmjob01	<u>Active</u>	<u>01</u>	ALLSERVERS

If you right-click a drill-down item in a column, the following options are available:

Name Column

Log and Statistics. Opens the Log and Statistics pane that enables you to view a log or statistics report.

The following options are available:

- Report Type.** Allows you to select either a log or statistics report.
- Report Range.** Allows you to select to view the last report, all the reports, or a report specified within a select time range.
- Completion Status.** Allows you to view all reports, all completed reports, all successfully completed reports, or all completed reports that contained errors.
- Run by.** Allows you to view reports for all flows, as well as those initiated by the scheduler or submit.
- Application or Flow Name.** Allows you to select a specific application or flow name for the report.

Status Column

- Manage Schedule.** Opens the Manage Schedule and E-Mail for procedure pane. For more information on this pane, see [How to Schedule a Repository Maintenance](#) on page 47.
- Scheduler Agents.** Runs a report for the user whose ID value is stored in the sched_scan_id of a Scheduler. The Scheduler must be active to run the Scheduler Agents report.
- Scheduled Events.** Opens a pane that shows all jobs scheduled to run during a specified time range.

The Scheduler must be active to run the Scheduler Events report.

This report shows all flows and procedures scheduled to run during the time period you selected. The Run Date and Run Time columns show their scheduled run times.

The User ID column shows the user ID used when running the flow or procedure.

Months Column

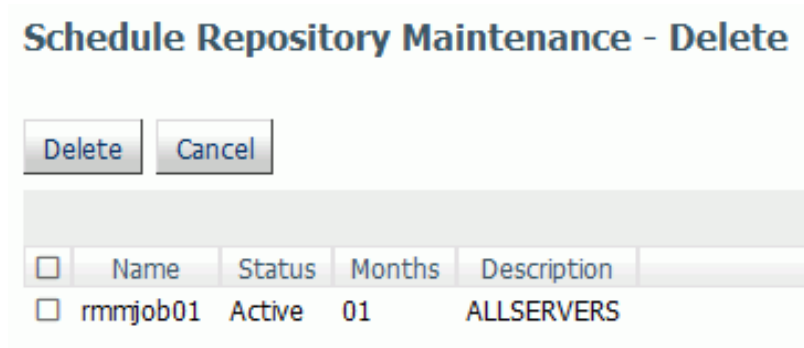
Older than. Opens the Resource Management Schedule Repository Maintenance - Change pane. Select the new time range of data that you want to delete and click Save.

Procedure: How to Delete Scheduled Jobs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Maintenance*, then point to *Schedule and E-Mail*, and click *Delete jobs*.

The Schedule Repository Maintenance - Delete pane opens.

4. Select the scheduled jobs you want to delete, as shown in the following image.



5. Click *Delete*.

Procedure: How to Submit a Maintenance Job


1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Maintenance*, and click *Submit job*.

The Resource Management Repository Maintenance pane opens, as shown in the following image.


Resource Management Repository Maintenance for FUENT-PC:8121

WARNING: The repository tables for this server should be backed up before proceeding.

All session and detail data in the date range for this server will be deleted

? Start Date:  2018/9/5


? End Date:  2018/9/6

 2018/9/7

12 ▾ 22 ▾

Submit

Note:

 The options for deleting data will vary based on how much data is in the repository.

4. Enter a Start Date and End Date for deletion of data through that range, or click the calendar icon to choose the dates.
5. Click *Submit*.

Note: Caution should be used when using this utility. Make sure you have a backup of the data before you proceed.

The Submit Repository Maintenance Job pane opens.

6. Enter a date and time for the maintenance job to be executed. Click *Submit*.

A Deferred Execution pane opens confirming your request.

Procedure: How to View a List of Scheduled Maintenance Jobs

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Maintenance*, and click *View jobs*.

The list of jobs opens, as shown in the following image.

Delete Selected		Refresh Selected									Q	▼
<input type="checkbox"/>	Defer ID	User ID	Status	App	Procedure	Report	Submitted Time	Scheduled Time	Days Left			
<input type="checkbox"/>	2018-01-30-15-31-07cmrpij000003_LOOPBACK	IBI\cf12026	Ready		RMUDEL		2018 January 31 16:08:39	---	29			
<input type="checkbox"/>	2018-02-01-10-29-29cmrpij000002_LOOPBACK	IBI\cf12026	Ready		RMUDEL		2018 February 01 11:40:26	---	29			

- To delete any of the scheduled maintenance jobs, select the check box next to the job you want to delete, and click *Delete Selected*.

Reconfiguring the Repository

This option allows you to convert the repository from either FOCUS to Relational or Relational to Relational. Before starting the reconfiguration process, you must create an adapter connection that points to the location of the new repository.

Procedure: How to Reconfigure the Repository

- Launch the Web Console.
- Click the *Resource Management* option on the sidebar.
- On the ribbon, click the *Repository* icon, point to *Reconfigure*, and click *Create new SQL repository*.

The Repository Reconfiguration pane opens, as shown in the following image.

Repository Reconfiguration: Create SQL Repository

Before starting, create a backup of the C:\ibi\srv99\wfs7709\catalog\rm directory

Reconfiguration to a new SQL Adapter consists of several steps. Step 1 is required. The rest of the steps are optional.

- Step 1. Create and initialize the new SQL repository.
- Step 2. Copy the Source repository data to the Target SQL repository.
- Step 3. Delete the temporary work files and Source repository.

Warning: If this process is interrupted Resource Management may be left in a state that will not allow data to be copied.

Step 1 - Create and initialize the new SQL data repository:

Select 'Yes' if you are authorized to create the SQL data repository.

Select 'DDL Only*' if you are not authorized to create the SQL data repository.

Select 'No' if the new configuration is to be added to an existing 'Many to One' repository.

*** Note: When the new data repository has been created, the Reconfigure process must be started again to finish the initialization of the new data repository**

? SQL Adapter	<input type="text" value="-- Select from list --"/>	
? Owner/Schema	<input type="text"/>	(case sensitive)
? Repository load type	<input type="text" value="Bulk load"/>	(Before selecting Bulk load confirm that the vendor's utility has been installed)
? Create Repository	<input type="text" value="-- Select --"/>	

- From the SQL Adapter drop-down menu, select the adapter and connection name for the new repository that was previously defined.
- Provide an Owner/Schema. By default, the tables are created for the specified owner/schema in the default database, and the owner/schema must be authorized to create tables in the database.

If the tables are to be created in an alternate database that the owner has rights to and the technique is supported by the RDBMS, the database name must be specified as well (for example, *database.owner*).

6. From the Create Repository drop-down menu, select *Yes*, unless the repository has already been created using the native Data Definition Language (DDL) of the data adapter or you are reconfiguring a many-to-one environment.

When *No* is selected, the configuration process will not issue DROP TABLE, DROP INDEX, CREATE FILE, and CREATE INDEX commands. Because the configuration process will always initialize the repository with default data and issue the default GRANTS, the repository must exist when the commands are executed or the configuration will fail.

If you are not authorized to create the SQL data repository, select *DDL Only*.

7. Click *Next*.

The Copy Data pane opens.

8. From the Copy data drop-down menu, select *Yes*.

Once you select *Yes*, you can select a date range to copy data in smaller portions, instead of all at once, using the calendar fields provided.

9. Click *Copy data*.

Once the data has been copied, a Deferred Execution pane opens showing the request.

***Procedure:* How to Submit a Reconfiguration Job**

Once you have created and initialized a SQL data repository, you can resubmit a copy job. This is done if there were any errors during the initial copy job, if the job did not run to completion, or if you want to select a different data range to copy. To submit a new job:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Reconfigure*, and click *Copy Data*.

The Repository Reconfiguration: Copy Data pane opens.

4. Specify the date range for the repository copy job.
5. Click *Copy data*.

Once the data has been copied, a Deferred Execution pane opens showing the request.

Procedure: How to View a Reconfiguration Job

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Reconfigure*, and click *View jobs*.

The Deferred List pane opens.

If you right-click an item in the Defer ID column, the following option is available:

- Delete.** Deletes the selected item.

Procedure: How to Delete Work Files

Once new repository and support files are created, you can delete temporary and old repository work files. To delete these files:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Repository* icon, point to *Reconfigure*, and click *Delete work files*.

The Repository Reconfiguration: Delete work files pane opens.

4. Click *Delete*.
5. Click *OK* to confirm the deletion.

The Repository Statistics pane opens, as shown in the following image.

Repository Statistics						
Note: The server you are connected to is listed in BOLD text						
Server Name	Monitoring	Earliest Date	Latest Date	Number of days	Total Sessions	Average Sessions
FUENT-PC:8121	Enabled	03/05/2018	03/05/2018	1	8	8.00

Migration

For more information on how to migrate Resource Management, see [Migrating the Resource Management Repository](#) on page 88.

Configuration Maintenance

The Configuration command has the following options:

- Properties
- Global Governing
- Configure Application Path
- Change License Key
- Delete Configuration

These give the administrator the option to see current configuration settings, as well as change the viewable application directories, change a license key, and remove the current Resource Management configuration.

Resource Management Properties

Setting monitor properties enables you to select global settings and the component levels to be monitored. Components are split into two categories, Procedure and Command. The information that Resource Governor monitors is stored in the Usage Monitoring tables. The default setting for Global Data Request monitoring is Off and the default setting for Component Level is Command. Session records are always captured.

Procedure: How to View Resource Management Properties

To view Resource Management Properties:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, point to *Properties*, and click *General*.

The General Properties pane opens.

The Resource Management Properties pane shows the following fields:

Configured

The date that Resource Management was configured or reconfigured.

Configuration Type

The product that is currently configured to Resource Management. Depending on which license key was used, the configuration type will be Resource Analyzer, Resource Governor, or Resource Analyzer Resource Governor.

Repository Type

The repository that is being used.

Data Adapter Connection

The data adapter connection of the current configuration.

Owner/schema

The owner/schema name for the relational adapter that Resource Governor will use to create the repository tables.

Global Data Request monitoring

Indicates whether Data Request Monitoring is on or off. When Global Data Request Monitoring is enabled, or set to *On*, Resource Governor collects information on all SQL SELECT and FOCUS TABLE, TABLEF, MATCH, GRAPH, MODIFY, and MAINTAIN requests. The default value is *Off*. Monitoring will be done by individually selected data sources.

Procedure Request Monitoring

Indicates the level of Procedure Request Monitoring.

WebFOCUS System Procedure Monitoring

Determines whether WebFOCUS internal system procedures will be monitored. The default setting is *Off*.

Data Sources monitored

The number of application data sources and relational pass-thru data sources that have been enabled for monitoring. This will show *Global monitor enabled* if Data Request Monitoring is on. See the Global Settings section below for more information.

Data Sources not monitored

The number of application data sources and relational pass-thru data sources that have been disabled by setting the monitoring setting to OFF.

Excluded Procedure count

The number of procedures that have been excluded from monitoring.

Last Maintenance Date

The date the repository was modified using the Repository Maintenance pane.

Last Repository Update

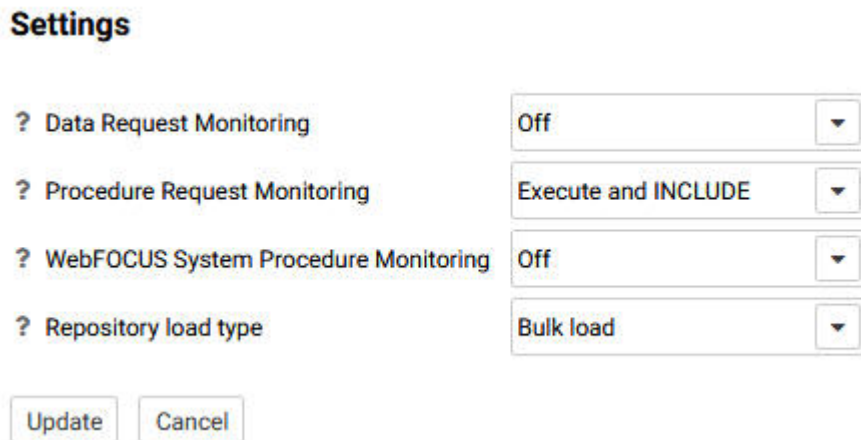
The last time the repository updated to include data from a log file.

Procedure: How to Set Monitor Properties

To set Resource Management Properties:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, point to *Properties*, and click *Settings*.

The Settings pane opens, as shown in the following image.



The Settings pane has the following options:

Data Request Monitoring

Indicates whether Data Request Monitoring is on or off. When Global Data Request Monitoring is enabled, or set to *On*, Resource Governor collects information on all SQL SELECT and FOCUS TABLE, TABLEF, MATCH, GRAPH, MODIFY, and MAINTAIN requests. The default value is *Off*. Monitoring will be done by individually selected data sources.

Procedure Request Monitoring

Indicates the level of Procedure Request Monitoring. The following options are available:

- Execute and INCLUDE.** All executed procedures and any procedure that used - INCLUDE will be monitored. This is the default value.
- Execute only.** Any procedure that is executed using EX or EXEC will be monitored.

- Primary Execute only.** The only procedure monitored will be the primary procedure. Any procedure executed from within the primary (-INCLUDE or EX) will not be monitored.

WebFOCUS System Procedure Monitoring

Determines whether WebFOCUS internal system procedures will be monitored. The default setting is *Off*.

Repository load type

Indicates the load method that will be used to add log data to the data repository. The following options are available:

- Bulk load.** The bulk load utility for the selected RDBMS server, if available, will be used to load the data into a temporary table on the RDBMS server using the connection adapter that was selected for the repository. This is the default value.
- Table load.** The TABLE command will be used to load the data into a temporary table on the RDBMS server using the connection adapter that was selected for the repository.

After the data is loaded into the temporary table using either load method, the MERGE command is then used to merge the data from the temporary table into the data repository.

4. Select the monitoring properties desired.
5. Click *Update*.

***Procedure:* How to Set Repository Properties**

The amount of data stored in the repository can be controlled by selecting the record types logged for each connected session. Collecting all of the records can give a complete overview of what each connected user is doing and is a good first step in determining what resources are being consumed and which procedures, data sources, and fields are being used and by whom. But collecting everything can stress the data repository, whether a FOCUS or Relational Database Management system is being used. Some of the tables are more geared towards the Resource Governor feature and others are used just for the Resource Analyzer feature. The Functions and WHERE/IF records are mainly used by the Resource Governor, but there are several reports in Resource Analyzer that will report on that data as well. From the Web Console Resource Management tab you have the option to add or remove any record type from the logging process using the Configuration/Properties/Repository menu in the Web Console ribbon or the Resource Management node on the tree.

To set Resource Management Repository Properties:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, point to *Properties*, and then click *Repository*.

The Repository Properties pane opens. The pane, as shown in the image below, shows the default record types collected. You can enable or disable any items on the pane to control when they are to be collected. Some records are required for others to be collected, so by selecting one record, it may cause dependent records to be selected automatically.

Repository Properties

^ Procedure requests

Procedure

Procedure parameters

^ Commands

SQL

TABLE/TABLEF

GRAPH

MATCH

APP

Command statements

JOIN

MODIFY

MAINTAIN

^ Command details

Data Sources

Fields

Fields used by Reference

Functions

WHERE/IF

^ Audit records

Translations

Messages

Access Control

Update

Cancel

Reset to Default

By default all procedures (focexec or SQL Stored Procedure) are monitored unless the Procedure requests log record has been disabled or the procedure has been removed from monitoring using the Remove Procedures from Monitor pane. By default, no data requests are collected unless one or more Data Requests log record types are selected and either Global Monitoring is enabled or at least one data source has been enabled for monitoring using the Add Data Sources to Monitor pane for each application folder. Also, the application must have Monitoring enabled in the Application properties pane.

Note: The All Data option in Release 7.6 also collects SORT and Functions. In Release 7.7, SORT has been included in Fields, and Functions has its own dialog box.

The Commands option in Release 7.7 also collects information about remote requests from a HUB-SUB configuration and, if active, Resource Governor records.

Resource Management will always collect session records and this feature cannot be turned off.

By default, command records (for example, a TABLE command being issued) will also be collected for a data resource that is not being monitored.

The following table shows the relationship between the Repository Properties and the log record types, as well as the corresponding repository entries.

Repository Property	Record Type	Repository Location	Comments
Procedure Requests			
Procedure	RLFX	SMRPCS	Each Procedure (for example, FOCEXEC or SQL Stored Procedure) executed will cause a RLFX record to be collected.
	RLFW	SMRPCWF	The RLFW record will only be collected if the session was started by a connection from the WebFOCUS Client or WebFOCUS self-service application. The RLFW record is added to the SMRPCWF table or segment. Only one RLFW record will be collected for the primary procedure, but several RLFX may be collected for a session.

Repository Property	Record Type	Repository Location	Comments
Procedure statements and parameters	RLFC	SMRPCREQUESTS (table) SMRPCSCMD (segment)	Contains the EX statement and the parameters or the -INCLUDE record. The actual command is executed. The RLFC records are collected if enabled and the Procedures are being collected.
Data Requests			
	RLCD	SMQUERY	These settings allow for the customization of command data to be collected. The RLCD record is collected and added to the SMQUERY table or segment. Each data request (for example, TABLE FILE... or SQL SELECT...) executed will cause an RLCD record to be collected. The RLCD can generate two additional child records, RLDG (Governor) and RLDE (request using SUFFIX=EDA data source). These two records cannot be disabled and will always be collected if generated.
	RLDG	SMGOVERN (table) SMGOVEND (segment)	
	RLDE	SMREMOTES (table) SMRMTS (segment)	
Data Request Details			
Request Statements	RLCC	SMRPCREQUESTS (table) SMCMDLN (segment)	Contains the complete request. For example, TABLE FILE CAR PRINT CAR END.
Functions	RLDU	SMFUNCTIONS (table) SMFNCTNS (segment)	Data about aggregate functions used on fields. There is one record for each function used containing the aggregation function used and the field it was used on.

Repository Property	Record Type	Repository Location	Comments
WHERE/IF	RLDR	SMRELATIONS (table) SMRELTNS (segment)	Data on the Relational tests used in a request, such as the operator, literals, or date and time values used. There is one record for each WHERE/IF/AND/OR and each record contains the right and/or left data source and field used. It also has the literal, if one was used on the left or right-hand side of the clause, and other information pertinent to a relational test clause.
Data Sources	RLDS	SMFROMS	The Master File Description/Synonym name or RDBMS table or tables used in a request. There is one record for each data source used in a request. This includes cross-referenced data sources or data sources joined using the JOIN or COMBINE command.
Field	RLDF	SMCOLUMNS (table) SMCOLMNS (segment)	The fields used, the type of field, and whether it was used in a sort or group in a request. There is one record for each field selected to return data in the request or in a BY or ORDER BY/GROUP BY statement. This record is not collected for fields used in functions or relational tests.
Fields used by Reference	RLDF	SMCOLUMNS (table) SMCOLMNS (segment)	Information about any DEFINE fields used or referenced in a request. The Fields used by Reference are indicated when the SMDEFINE field is set to 5 (Referenced). Referenced fields are fields that are used in a DEFINE in the MFD, DEFINE FILE, or COMPUTE statement to create a virtual field, and the virtual field is used in the request.

Audit Records

Repository Property	Record Type	Repository Location	Comments
Translations	RLAU	SMAUDIT	Contains the translated request if translation is required. If a WebFOCUS TABLE request is for a relational data source, the translated SQL, which is sent to the RDBMS, will be monitored. The RLAU records are collected if enabled and the data request statements are being collected.
Messages	RLAU	SMAUDIT	Messages that are issued during the processing of a request will be monitored.
Access Control	RLAU	SMAUDIT	Information about the changes to server configuration files, applications, and files. The data shows the timestamp, the user ID of who made the change, and the TCP address the change originated.

The following table shows the comparisons between the Monitoring settings in 7.6, and earlier releases, and the Command Log Records settings in Release 7.7.

Release 7.6	Release 7.7
Monitoring Settings	Command Log Records
All Data	Select all
Query, Froms, Relations, and Columns	Commands, Request, Data Sources, WHERE/IF, and Fields
Query, Froms, and Relations	Commands, Request, Data Sources, and WHERE/IF
Query and Froms	Commands, Request, and Data Sources
Query only	Commands and Requests

4. Click *Update*.

Global Governing

For more information on global governing, see [Working with Knowledge Bases](#) on page 115.

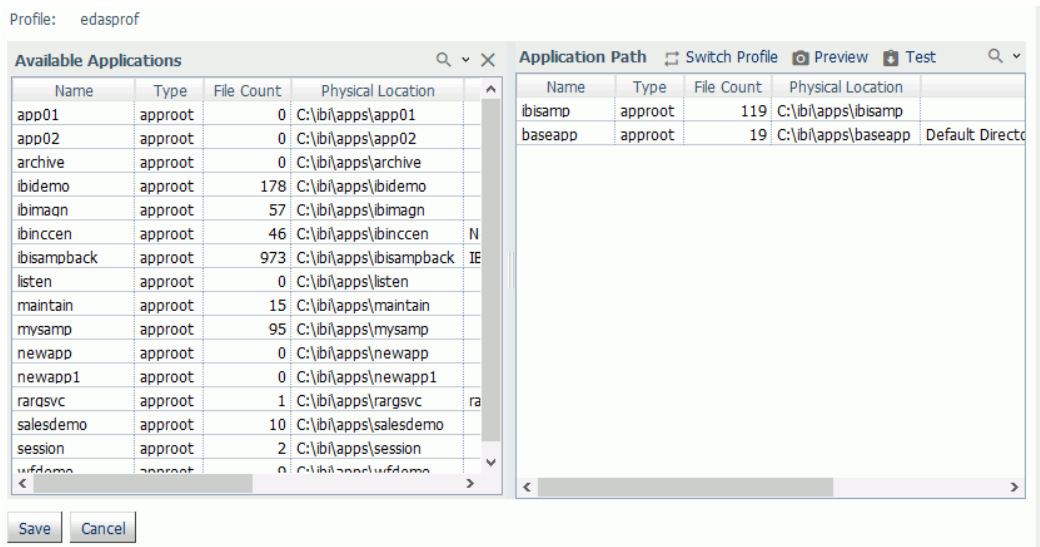
Configuring Application Paths

You can configure the application path to set the viewable directories for Resource Governor.

Procedure: How to Configure an Application Path

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, and click *Configure Application Path*.

The Application Path Configuration pane opens, as shown in the following image.



4. To add or remove a directory from the application path, right-click the application and click *Add to Path* or *Remove from Path*, respectively.
5. Click Save when finished.

Refresh the navigation pane to see the newly added or removed directory.

Migrating Resource Management

For more information on how to migrate Resource Management, see [Migrating the Resource Management Repository](#) on page 88.

Working With License Keys

A license key is required to run Resource Governor. You can change the license key you are using from within Resource Management at any time.

Procedure: **How to Change the License Key**

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, and click *Change License Key*.

The Resource Management Configuration pane opens showing your current license.

4. Enter the new license into the License field and click *Save*.

Removing Resource Management

You can remove Resource Management from the server from the Resource Management Configuration pane. This is required if you want to change the repository type from FOCUS to relational, or relational to FOCUS.

Procedure: **How to Remove Resource Management**

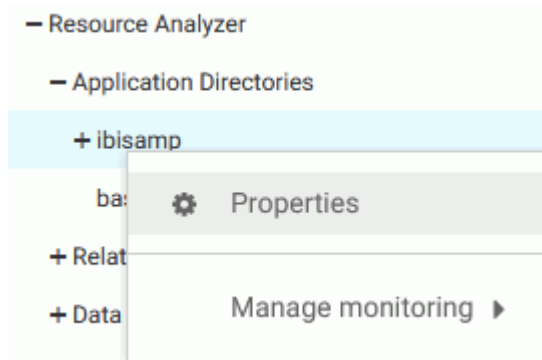
1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, click the *Configuration* icon, and click *Delete Configuration*.

The Delete Configuration pane opens.

4. Click *Delete* to remove Resource Management from the server. All system files and the repository will be removed from the server.

Working With Application Directories

When you right-click an application directory, a context menu is available, as shown in the following image.



The following sections provide detailed explanations for each option in this context menu.

Monitoring at the Application Level

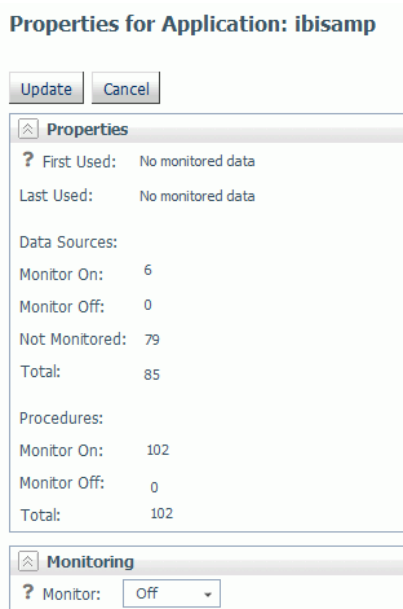
Monitoring Data Sources can be set at the application level to monitor all data sources or at the individual level.

Note: Monitoring at the individual data source level will take precedence over the application level. For more information, see [Monitoring Individual Data Sources](#) on page 75.

Procedure: How to View Application Name Properties

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Application Directories* folder in the tree.
4. To view the data source properties, right-click an application directory and click *Properties*.

The Properties for Application pane opens, as shown in the following image.



The Application pane has the following fields/options:

First Used

Is the date the application was first accessed.

Last Used

Is the date the application was last accessed.

Data Sources

Monitor On

The number of data sources with monitoring on.

Monitor Off

The number of data sources with monitoring off.

Not Monitored

The number of data sources not being monitored.

Total

The total number of data sources in the selected application directory.

Procedures

Monitor On

The number of procedures with monitoring on.

Monitor Off

The number of procedures with monitoring off.

Total

The total number of procedures in the selected application directory.

Monitor Status

For more information on setting monitor status, see [How to Set Monitor Status at the Application Level](#) on page 74.

Note: The samples application directory *ibisamp*, created at product installation, has monitoring turned off, by default.

Procedure: How to Set Monitor Status at the Application Level

Data sources in an application can have monitoring turned on or off at the application level. When *On* is selected, the Usage Monitor will collect information about all requests for all data sources in the application. When *Off* is selected, no monitored data for data sources in the application is collected.

If Application monitoring has never been used, the Monitor Status will be *Not Set*. This is the default value.

To set monitor status for an application:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand an application directory in the tree.
4. To view the Application properties, right-click an application name and click the *Properties* option from the context menu.

The Properties for Application pane opens.

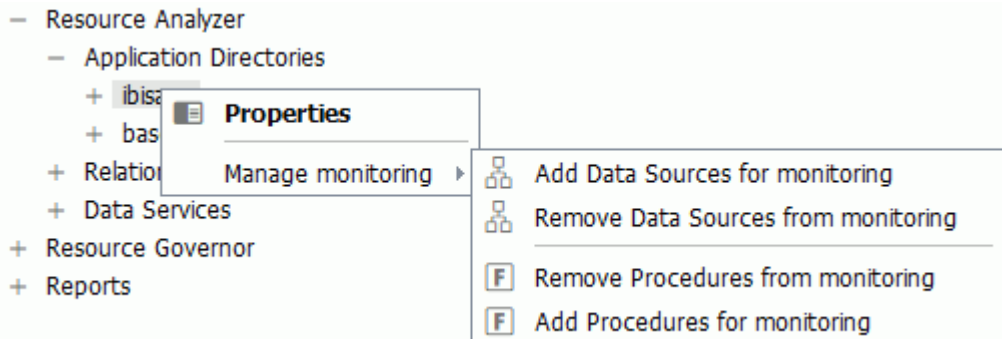
5. Select the monitoring status you want to use from the Monitor Status drop-down menu. The options are:

- Not Set.** This is the same as *Off* and is the default. The status will not be shown in the navigation tree when viewing data sources.

- ❑ **On.** Turns monitoring for the application on and can be seen in the navigation tree when viewing the Application list.
- ❑ **Off.** Turns monitoring for the application off and can be seen in the navigation tree when viewing the Application list.

Monitoring Individual Data Sources

When you right-click an application directory and point to *Manage monitoring*, a context menu is available, as shown in the following image.

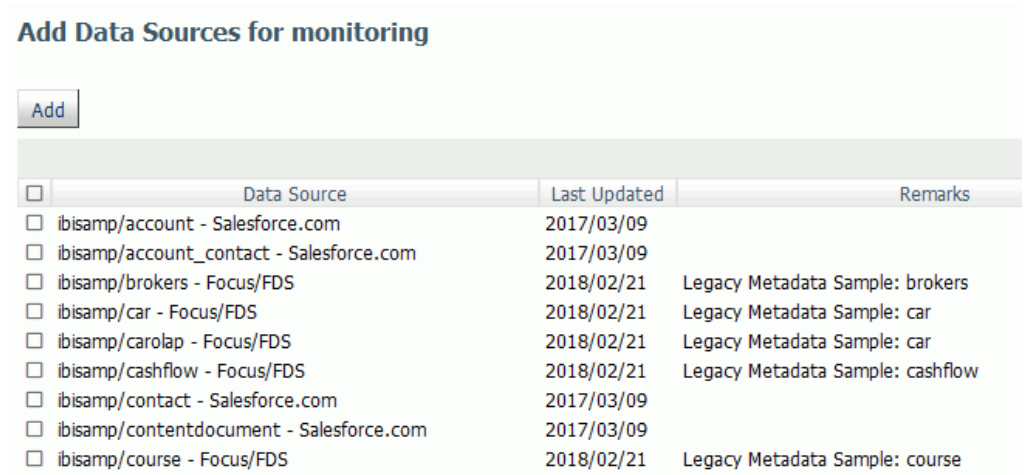


Procedure: How to Monitor Individual Data Sources

Data sources must be monitored to provide monitor data, which can then be used to analyze performance.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click an application directory in the navigation pane, point to *Manage monitoring*, and click *Add Data Sources for monitoring*.

The Add Data Sources for monitoring pane opens to display all synonyms in the selected application directory, as shown in the following image.



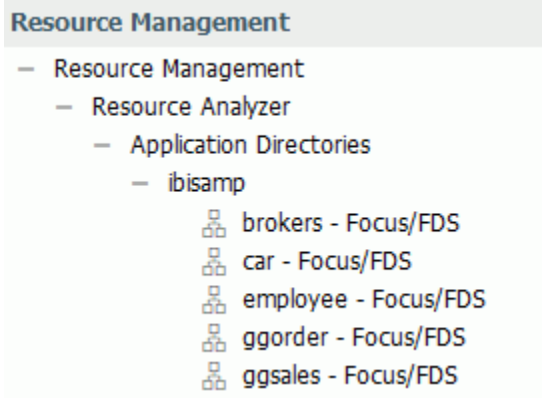
4. Select the check boxes for the data sources you want to monitor, or select the Select All check box next to the Data Source Name column to select all of the data sources listed.
5. Click *Add*.

The Add Data Sources for monitoring - Status pane opens displaying the selected data sources with monitoring enabled, as shown in the following image.



6. Click the *Close* button.

The data sources you selected now appear in the tree, as shown in the following image.



Procedure: How to Remove a Monitored Data Source From Being Monitored

You can remove a data source from being monitored as follows.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click an application directory in the navigation pane, point to *Manage monitoring*, and click *Remove Data Sources from monitoring*.

A list of available data sources opens.

4. Select the data sources that you want to remove from monitoring or select the *Select All* check box next to the *Data Source Name* column, and click *Remove*.

The *Remove Data Sources from monitoring - Status* pane opens, confirming that the selected data source was removed.

5. Click *Close*.

Procedure: How to Remove a Procedure From Being Monitored

By default, all application data source procedures are monitored. You can remove a procedure from being monitored as follows.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click an application directory in the navigation pane, point to *Manage monitoring*, and click *Remove Procedures from monitoring*.

A list of available procedures opens.

4. Select the procedures that you want to remove from monitoring or select the Select All check box next to the Procedure Name column, and click *Remove*.

The Remove Procedures from monitoring - Status pane opens, confirming that the selected procedure was removed.


5. Click *Close*.

Procedure: How to Monitor Procedures

To have the ability to add procedures, they will need to be removed from monitoring first. For more details, see [How to Remove a Procedure From Being Monitored](#) on page 77.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click an application directory in the navigation pane, point to *Manage monitoring*, and click *Add Procedures for monitoring*.

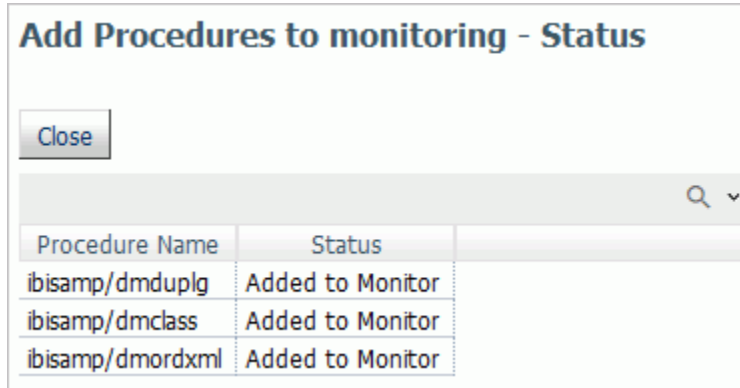
The Add Procedures to monitoring pane opens to display all procedures in the selected application directory, as shown in the following image.



4. Select the check boxes for the procedures you want to monitor, or select the Select All check box next to the Procedure Name column to select all of the procedures listed.

5. Click *Add*.

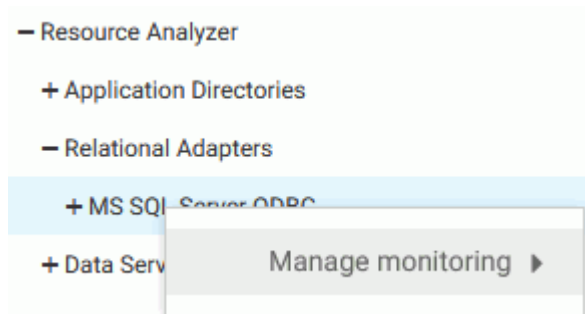
The Add Procedures to monitoring - Status pane opens displaying the selected procedures with monitoring enabled, as shown in the following image.



6. Click the *Close* button.

Monitoring Relational Adapters

When you right-click a relational adapter folder, a context menu appears, as shown in the following image.



The following sections provide detailed explanations for each option in this context menu.

Relational Adapter Properties

Procedure: How to View Relational Adapter Properties

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.

3. Expand the *Relational Adapters* folder.
4. Right-click a relational adapter directory in the tree and select *Properties*.
5. Select a connection name from the Connection Name drop-down menu.

Note: If there is only one connection name available, the properties for that connection will be shown immediately.

The Properties for Relational Adapter pane opens for the selected connection, as shown in the following image.

Properties for Relational Adapter - MS SQL Server ODBC	
Connection Name	CON01 ▾
First Used:	Never
Last Used:	Never
Data Sources:	
Monitored:	0
Monitor Off:	0
Not Monitored:	497
Total:	497
Procedures:	
Monitor Off:	0
Total:	1513

The Properties for Relational Adapter pane has the following fields/options:

Connection Name

Allows you to select the connection name from the drop-down menu.

First Used

Is the date the relational adapter was first accessed.

Last Used

Is the date the relational adapter was last accessed.

Data Sources

Monitored

The number of data sources with monitoring on.

Monitor Off

The number of data sources with monitoring off.

Not Monitored

The number of data sources not being monitored.

Total

The total number of data sources in the selected application directory.

Procedures

Monitor Off

The number of procedures with monitoring off.

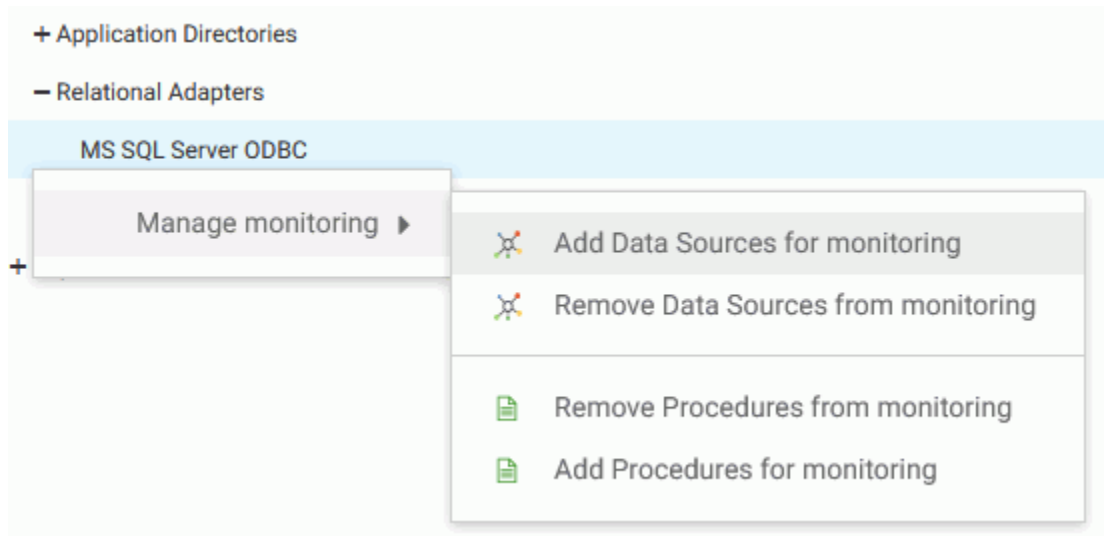
Total

The total number of procedures in the selected application directory.

Monitoring Relational Data Adapter Sources

When you right-click a relational adapter and point to *Manage monitoring*, a context menu is available, as shown in the following image.

In the example shown, MS SQL Server ODBC is the adapter selected for relational source monitoring.



Procedure: How to Monitor Relational Data Adapter Sources

Monitoring relational data adapter sources means monitoring SQL direct passthru requests to relational sources that do not use synonyms. Relational sources must be monitored to provide monitor data, which can then be used to analyze performance or govern usage.

Note: The following procedure can only be done if a relational adapter has been configured.

To monitor relational sources:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Relational Adapters* folder in the tree, right-click the appropriate data adapter, point to *Manage monitoring*, and click *Add Data Sources for monitoring*.
4. Select the connection name from the drop-down menu.

Note: If there is only one connection name available, the pane will be shown immediately.

The Add Data Sources for monitoring pane opens, as shown in the following image.

Add Data Sources for monitoring for - MS SQL Server ODBC/CON01

Connection Name

Restrict search to:

Tables

Views

Filter by owner/schema and object name

5. If required, select the *Filter* check box and provide filter information in the appropriate fields. Click *Next*.

The Add Data Sources for monitoring pane opens for the connection, as shown in the following image.

Add Data Sources for monitoring for - MS SQL Server ODBC/CON01

<input type="checkbox"/>	Data Source	Monitor As	Type	Comment
<input type="checkbox"/>	dbo.wrd_wf_retail_age	Owner.Table	▼ TABLE	
<input type="checkbox"/>	dbo.wrd_wf_retail_currency	Owner.Table	▼ TABLE	
<input type="checkbox"/>	dbo.wrd_wf_retail_customer	Owner.Table	▼ TABLE	
<input type="checkbox"/>	dbo.wrd_wf_retail_discount	Owner.Table	▼ TABLE	
<input type="checkbox"/>	dbo.wrd_wf_retail_education	Owner.Table	▼ TABLE	
<input type="checkbox"/>	dbo.wrd_wf_retail_employee	Owner.Table	▼ TABLE	

Note: Any data source with special characters will be quoted using the RDBMS quoted identifier character. The data source will be stored as a two part name, and monitored using upper case characters, with the exception of any quoted special characters.

6. Select the check boxes for the data sources you want to monitor, or select the Select All check box next to the Data Source column to select all of the data sources listed.

7. From the Monitor As column, select one of the following options from the drop-down menu:
 - Table.** Monitor the data source using the one part name.
 - Owner.Table.** Monitor the data source using the two part name. This is the default value.
 - Both.** Monitor the data source using both the one part and two part names.
8. Click *Add*.

The Add Data Sources pane reopens displaying the selected data sources with monitoring enabled.
9. Click *OK* to close the confirmation dialog box.
10. Click *Close*. The relational data sources you selected now appear in the tree under the adapter/connection folder.

Procedure: How to Remove a Data Source for a Relational Adapter From Being Monitored

You can remove a data source from being monitored as follows.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click a relational adapter directory, point to *Manage monitoring*, and click *Remove Data Sources from monitoring*.

The *Remove Data Sources from monitoring* pane opens.

4. Select the connection name from the drop-down menu.
5. If required, select the *Filter* check box and provide filter information in the appropriate fields. Click *Next*.

The Remove Data Sources pane opens.

6. Select the check boxes for the data sources you want to remove, or select the Select All check box next to the Data Source column to select all of the data sources listed. Click *Remove*.

The Remove Data Sources pane reopens, displaying the data sources that were removed.

7. Click *OK* to close the confirmation dialog box.
8. Click *Close*.

Procedure: How to Remove a Procedure for a Relational Adapter From Being Monitored

By default, all relational procedures are monitored. You can remove a procedure from being monitored as follows.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click a relational adapter directory, point to *Manage monitoring*, and click *Remove Procedures from monitoring*.

The Remove Stored Procedures from monitoring pane opens, as shown in the following image.

Remove Stored Procedures from monitoring for - MS SQL Server ODBC

Connection Name

Filter by owner/schema and object name

4. Select the connection name from the drop-down menu.
5. If required, select the *Filter* check box and provide filter information in the appropriate fields. Click *Next*.

The Remove Stored Procedures from monitoring pane opens, showing the list of available procedures.

6. Select the procedures you want to remove from monitoring, and click *Remove*.

The Remove Procedures from Monitor - Status pane opens, confirming that the selected procedure was removed.

7. Click *OK* to close the confirmation dialog box.
8. Click *Close*.

Procedure: How to Monitor Procedures

To have the ability to add procedures, they will need to be removed from monitoring first. For more details, see [How to Remove a Procedure for a Relational Adapter From Being Monitored](#) on page 85.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.


3. Right-click a relational adapter directory in the navigation pane, point to *Manage monitoring*, and click *Add Procedures for monitoring*.

The Add Stored Procedures to monitoring pane opens.

4. Select the connection name from the drop-down menu.
5. If required, select the *Filter* check box and provide filter information in the appropriate fields. Click *Next*.
6. Select the check boxes for the procedures you want to monitor, or select the *Select All* check box next to the *Stored Procedure* column to select all of the listed procedures.
7. Click *Add*.

The Add Stored Procedure for monitoring pane opens displaying the selected procedures with monitoring enabled, as shown in the following image.

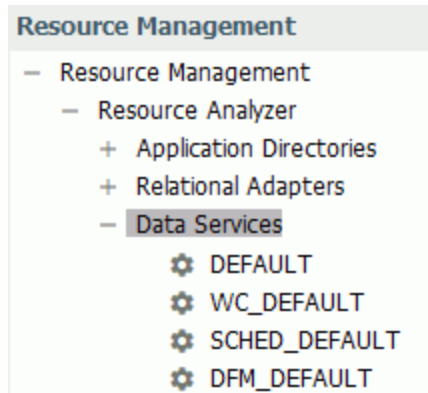
Add Stored Procedure for monitoring for - MS SQL Server ODBC/CON01



8. Click *OK* to close the confirmation dialog box.
9. Click the *Close* button.

Monitoring at the Data Service Level

When you expand the *Data Services* folder in the navigation pane, the data services available on the server are listed. An example is shown in the following image.




Procedure: How to Disable Monitoring on Data Services

All defined Data Services are displayed in the Resource Management navigation pane.

By default, all Data Services are monitored. To disable monitoring:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Data Services* folder.
4. Right-click the Data Service you wish to disable and click *Properties*.

The Data Services Properties pane opens, as shown in the following image.



5. Select *OFF* from the Monitor drop-down menu.
6. Click *Update*.
A confirmation dialog box opens.
7. Click *OK*.

The Data Services node in the navigation pane now shows the Data Service with monitoring disabled, as shown in the following image.



Procedure: How to Enable Monitoring on Data Services

All defined Data Services are displayed in the Resource Management navigation pane.

To enable monitoring on a Data Service that was previously disabled:

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Data Services* folder.
4. Right-click the disabled Data Service you wish to enable and click *Properties*.

The Data Services Properties pane opens.

5. Select *ON* from the Monitor drop-down menu.
6. Click *Update*.

A confirmation dialog box opens.

7. Click *OK*.

The Data Services node in the navigation pane now shows the Data Service with monitoring enabled.

Migrating the Resource Management Repository

If you want to access monitor data from a previous release when upgrading Resource Governor, you must migrate the previous Resource Management repository contents into the new release repository.

Requirements for Migrating the Resource Management Repository

The following conditions are required to migrate the Resource Management repository.

- ❑ Monitoring must be turned off on both the host and target server before running the migration procedure. Also, the host log file must be archived to the repository before it is migrated to the new repository.

To turn off monitoring, use the server Web Console to Disable Monitoring & Governing. If you turn monitoring back on and continue to monitor in the old release, you must re-run the migration in order to migrate the newly monitored data.

- ❑ When migrating monitor data stored in FOCUS, the tables must be available during migration. No actively running sink machine (SU) process is recommended or required.
- ❑ When migrating monitor data stored in a relational repository, an appropriate database adapter must be configured with a connection pointing to the data from the prior release. If the new relational repository is to be stored in the same location as the old data, appropriate values for owner, creator, and so on must be used to differentiate the two sets of data.
- ❑ For both relational and FOCUS migration, the repository tables in the old release locations remain where they were and can be used as backup.

Note: When you run the migration more than once, which is necessary if you continued monitoring in the old release after migration and want the newly monitored data to be migrated, you will get a duplicate record error message. The duplicates are ignored and the new records are added.

Migrating a Resource Management Repository for a New Server Install

You can migrate a Resource Management repository on UNIX, Windows, UNIX System Services, and IBM i for a new server install. For these platforms, the migration procedure is executed in a Web Console session.

Note: IBM i was formerly known as i5/OS.

Procedure: How to Migrate a Resource Management Repository for a New Server Install

To start the migration job:

1. Click the *Resource Management* option in the sidebar.
2. On the ribbon, in the Repository group, click *Migration* and click *Migrate old release*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and then click *Migrate old release*.

Note: You can also perform this task by clicking the *Workspace* option in the sidebar, right-clicking *Workspace* in the navigation pane, and clicking *Migrate*.

3. The Repository Migration page opens, as shown in the following image.



The screenshot shows a dialog box titled "Repository Migration". It contains two input fields with question mark icons to their left. The first field is labeled "Release to migrate data from" and has a dropdown menu with the text "- Old Release -" and a downward arrow. The second field is labeled "EDACONF directory of the old release" and is an empty text box with a three-dot menu icon to its right. Below these fields is a "Cancel" button.

4. From the drop-down menu, select the release number from which you are migrating. Options include 7703, 7704, 7705, 7706, 7707, and 7708.
 5. For all releases, you must enter the path to the EDACONF directory of the release being migrated. Examples are C:\ibi\svr77\wfs for Windows and /home1/ibi/svr77/wfs for UNIX.
 6. Enter the path to the EDAHOME directory of the release being migrated. Examples are C:\ibi\svr77\home for Windows and /home1/ibi/svr77/home for UNIX.
- Note:** This is only required if the old release is Version 7.7.03.
7. Click *Next*.

The Resource Management Migration: Submit Migration Job pane opens, as shown in the following image.

Submit Migration Job

Release to be migrated: **7706**

Old EDACONF: **C:\ibi\srv77\wfs7706**

Repository data range: **03/29/2018 - 03/29/2018**

System Database

? Migrate

Data Repository

? Migrate

- From the System Database Migrate drop-down menu, select Yes if you want to migrate the old release system database. If you select Yes, additional options become available, as shown in the following image.

Submit Migration Job

Release to be migrated: **7706**
Old EDACONF: **C:\ibi\srv77\wfs7706**
Repository data range: **03/29/2018 - 03/29/2018**

System Database

? Migrate

Yes should only be selected once, or rejected record counts will be noted in in the subsequent migration jobs output

Configuration data

Selecting *Yes* will replace existing settings that may have already been set

? Log record preferences

? Logging Properties

Data Repository

? Migrate

The additional Configuration data options are:

Log record preferences

Select Yes to migrate the old release log record preference settings to the system repository of the current server. Migrating this data will cause any existing log record preference data to be replaced with the data from the old release.

Logging Properties

Select Yes to migrate the old release logging property settings to the system repository of the current server. Migrating this data will cause any existing logging properties data to be replaced with the data from the old release.

- From the Data Repository Migrate drop-down menu, select Yes if the previous release data repository should be migrated. If Yes is selected, SMCNTRL, SMPRMTRS, and SMPRL data will be migrated. The default value is No.

If you select Yes, additional options become available, as shown in the following image.

Data Repository

? Migrate

No data has been migrated at this time

The migration job will be submitted using the selected date range below when 'Submit Job' is clicked.

? Start Date

? End Date

Migration jobs using different date ranges can be submitted again by clicking the 'Migration/Submit Migration job' menu item. Do not submit another job until the current job completes and the results have been reviewed

The additional options let you migrate data in smaller portions, instead of all at once. The available options are:

Start Date

The starting date of the data to be migrated. The default value is the earliest date in the old repository.

End Date

The ending date of the data to be migrated. The default value is the latest date in the old repository.

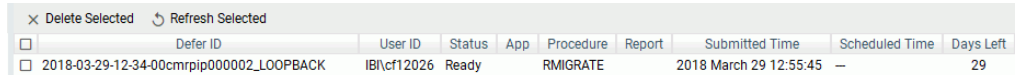
Note: Only custom BRL members will be migrated. The SMKNBNAME value in SMCNTRL will not be migrated, and any Govern and/or Advise values will be set to OFF. Any compiled rule files must be rebuilt after the migration is completed, and new Govern and/or Advise values must be set.

10. Accept the default dates, or enter a specific date that you want the migration to run.
11. Click *Submit Job*. A confirmation of the request appears, stating that the deferred execution request is in the queue.
12. On the ribbon, in the Repository group, click *Migration* and then click .

or

From the navigation pane, right-click , point to *Migration*, and click .

The Deferred List pane opens, displaying the list of deferred requests, and showing the request name that was displayed when the deferred migration job was scheduled. Examples of request names are listed in the Defer ID column, as seen in the following image.



<input type="checkbox"/>	Defer ID	User ID	Status	App	Procedure	Report	Submitted Time	Scheduled Time	Days Left
<input type="checkbox"/>	2018-03-29-12-34-00cmrrip000002_LOOPBACK	IB\cf12026	Ready		RMIGRATE	RMIGRATE	2018 March 29 12:55:45	—	29

The Status column shows the state of the deferred request, which can be Queued, Executing, or Ready. To see if the status has changed to Ready, either keep refreshing the page using the refresh options located at the top of the page, or redisplay the page at a later time. Once the status is Ready, right-click the request line and click *Get* to review the job output.

13. You can choose whether to remove the old configuration files or perform this action in the future.

On the ribbon, in the Repository group, click *Migration* and then click *Delete work files*.

or

From the navigation pane, right-click *Resource Management* point to *Migration*, and click *Delete work files*.

- a. Click *Delete* and then *OK* to confirm.

Note: Once you choose to remove the old configuration files, the option to migrate data will no longer be available.

Migrating a Resource Management Repository After a Server Refresh

You can migrate a Resource Management repository on UNIX, Windows, UNIX System Services, and IBM i after a server refresh. The migration information described below is only applicable if the server environment has been upgraded by performing the refresh option during the installation process. To refresh the server, follow the instructions in the Installation Guide for your platform.

Migration Usage Notes:

- The migration process must be applied to the original server that was refreshed. If a clone server was set up and the Resource Management repository was copied to this server for the purpose of testing, the migration process will not work. There are inter-dependencies in the Resource Analyzer repository data based on the machine name and port number used when first configured.

- The migration of the Resource Management repository on MVS PDS deployment after a refresh is not supported. A new server environment is required.

Once the refresh process has completed successfully, start the server and then the Web Console.


The migration is performed from the Web Console and consists of three phases:

- Migrate configuration file (Required).
- Migrate repository (Optional).
- Remove old files (Optional).

For these platforms, the migration procedure is executed on a Web Console session.

Note: IBM i was formerly known as i5/OS.

If you start the Web Console session with security ON, a message is displayed, as shown in the following image.



This message indicates that the structure of the Resource Management repository has changed. A migration process is required in order to update the existing repository so that new data columns can be archived to the repository. If the migration process is not run, monitor data will still be collected and archived but without the new data columns.

You can:

- Ignore the message contents by closing the message window. The repository will still have the old structure and the warning message will continue to be shown for a total of 10 times, or until the migration process has been performed.
- Start the migration process by clicking *Continue* and then clicking *Resource Management*.

The Resource Management Migration pane opens. The sequence of panes is slightly different depending on the repository type.

Note: If you start the Web Console session with security OFF, you will not see the message as described above. The Resource Management Migration pane will be immediately available after clicking *Resource Management*.

Procedure: How to Migrate a Repository After a Server Refresh

The following procedure details the steps for migrating a FOCUS or relational repository after a server refresh.

1. Click the *Resource Management* option in the sidebar.
2. **For a FOCUS Repository: Phase 1 - Migrate Configuration**, the pane opens in Phase 1.
 - a. Click *Reconfigure* to complete Phase 1.
3. **For a Relational Repository: Phase 1 - Migrate Configuration**, the pane opens in Phase 1, as shown in the following image.

Repository Migration

The server was upgraded to a new release and new features for Resource Management must be added to the repository and configuration files.

Before starting, make backups of the following configuration files:

- *rmprof.prf* is located in: C:\ibi\srv77\wfs7706d\etc
- *gktable.prf* is located in: C:\ibi\srv77\wfs7706d\etc
- All files prefixed with 'RML' located in: C:\ibi\srv77\wfs7706d\catalog\rm

The migration process consists of several steps. Step 1 is required. The rest of the steps are optional.

- Step 1. Upgrade configuration files and create the repositories with new features.
- Step 2. Migrate existing repository data to the new repositories.
- Step 3. Delete the temporary work files and old repository.

Step 1

Create Repository

Resource Management will perform the following actions when the 'Reconfigure' button is clicked.

- Rename existing RMPROF to *rmpr7706.prf*
- Rename existing GKTABLE to *gktb7706.prf*
- Create and initialize the new RMLSYS repository using the existing RMLSYS configuration values.
- Create and initialize the new RMLDB repository.

After step 1 has completed the new RMLDB and RMLSYS repositories will have been initialized and ready for archiving log data.

The options in the Create Repository drop-down menu are *Yes*, *No*, and *DDL Only*.

For Phase 1 - Create Repository Yes:

Select this option if you have DBA authority over the existing Resource Management repository tables. New repository tables will be created with a suffix of *_nnnn*, where *nnnn* reflects the release number. The process will continue at Phase 2.

- a. Select *Yes* from the Create Repositories drop-down menu. A list of actions that will be performed is displayed.
- b. Click *Next* to go to step 4 and begin Phase 2.

For Phase 1 - Create Repository DDL Only:

Select this option if you do not have DBA authority over the existing Resource Management repository tables. This option creates a file, `rmlldb.sql`, which contains RDBMS specific DDL.

- a. Select *DDL Only* from the Create Repositories drop-down menu.
- b. Click *Create DDL*. The instructions for the location of the `rmlldb.sql` file are displayed. Give this file to the DBA for processing.
- c. Click *Close*. You will be returned to the first migration screen. You must create the tables before continuing.

Note: Your DBA must create the new repository tables in the same location as the original 7.7 tables, and both sets of tables have to be accessible on the same adapter connection on the server. The new tables will have the release number appended to the name. The format will be `tablename_release`.

- d. Once the repository tables have been created, return to the Web Console and click *Resource Management*.
 - e. Go to step 4 to begin Phase 2.
4. **Phase 2 - Migrate data**, select the Migration data option from the drop-down menu, as shown in the following image.

Submit Migration Job

Step 1 - Migrate configuration actions:

The system repository has been created and initialized and all old system data has been migrated.

The data repository has been created and initialized and ready for data migration.

Step 2 - Migrate data:

No data has been migrated at this time

The date range of data in the source repository: **04/11/2018 - 04/11/2018**

? Migrate data

The repository migration can be submitted when you are ready by clicking the 'Migration/Submit Migration job' menu item.

The options are Yes or No.

Note:

- If you selected *DDL Only* in the previous step, click *Reconfigure*.
- When returning to this page after selecting *DDL Only*, the default value is *No*. Do not select *Yes* or an error will occur.

For Phase 2 - Migrate data Yes:

Select this option to copy data from the old repository to the new one. When this option is selected, additional options become available, as shown in the following image.

Step 2 - Migrate data:

No data has been migrated at this time

The date range of data in the source repository: 04/11/2018 - 04/11/2018

? Migrate data

The repository migration job will be submitted using the selected date range when the 'Run Job' button is clicked.

? Start Date

? End Date

The additional options let you migrate data in smaller portions, instead of all at once. The available options are:

Start Date

The starting date of the data to be migrated. The default value is the earliest date in the old repository.

End Date

The ending date of the data to be migrated. The default value is the latest date in the old repository.

- a. Click *Submit Job* to schedule the migration procedure. A confirmation of the request appears, stating that the deferred execution request is in the queue.
- b. Go to step 7.

For Phase 2 - Migrate data No:

Select this option if you do not want data to be copied from the old repository to the new one.

To run the migration process at a later date, select *Resource Management* on the sidebar and on the ribbon, in the *Repository* group, click *Migration*, and click *Submit job*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and click *Submit job*.

The Resource Management Migration: Submit Migration Job pane will open on Phase 2.

- a. Click *Exit* to continue.
 - b. Optionally, review step 5.
5. To view the Migration Job page if migration was selected, on the ribbon, in the Repository group, click *Migration*, and then click *View job*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and click *View job*.

The Deferred List pane opens.

6. Right-click on the job and click *Get* to see if there are any FOCUS errors.
7. **Phase 3 - Remove old configuration and data files**, you can choose whether to remove the old configuration files or perform this action in the future.

On the ribbon, in the Repository group, click *Migration*, and then click *Delete work files*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and click *Delete work files*.

- a. Click *Delete* and then *OK* to confirm.

Note: Once you choose to remove the old configuration files, the option to migrate data will no longer be available.

Migrating a Resource Management Repository on MVS

This option is only available when migrating data from a new server install of Release 7.7. Repository migration is not available from earlier releases.

Procedure: How to Migrate a Resource Management Repository on MVS

To start the migration job:

1. Click the *Resource Management* option in the sidebar.
2. On the ribbon, in the Repository group, click *Migration* and click *Migrate old release*.

or

From the navigation pane, right click *Resource Management*, point to *Migration*, and then click *Migrate old release*.

3. The Resource Management Migration page opens.
4. To migrate data from the release, use the drop-down menu to select the release number from which you are migrating. The options include 7700-7702, 7703, 7704, 7705, and 7706. Depending on the release selected, one or more of the following options will be available.

Note: For help in identifying the dataset names required on the next two panes, see the chart in Step 6.

GKTABLE dataset name

The dataset name of the Resource Management parameter file.

REPOSITORY dataset name

The dataset name of the FOCUS repository from the previous release.

RMLSYS dataset name

The dataset that contains Resource Management system information.

CONFIGURATION dataset name

For old release selections of 7703 or 7704, this is the dataset name that contains the Master File profile members. They are RMLPD773 (7703) or RMLPD774 (7704) and RMLPRFSY (both releases).

5. Click *Next* to continue.
6. The Submit Migration Job pane opens.

For all other releases, the same pane will show but without the Create 7.6 compatibility masters option available. The following options are available:

Migrate system data

Select Yes if the previous releases system data should be migrated. If Yes is selected, SMCNTRL, SMPRMTRS and SMPRL data will be migrated. Only custom BRL members will be migrated. The SMKNBNAME value in SMCNTRL will not be migrated and any Govern and/or Advise values will be set to OFF. Any compiled rule files must be rebuilt after the migration is completed, and new Govern and/or Advise values must be set. The default value is *No*.

For releases 7700 – 7702, 7703 and 7704, if Yes is selected, one or two additional input boxes are shown depending on the old release selection.

All of the names in the following chart require *QUALIF.servertype* as a prefix.

Where *servertype* can be FFS, WFS, or DM.

Dataset name	7704-7706
GKTABLE	CONF.PRF
REPOSITORY**	RMLDB.FOCUS
RMLSYS*	RMLSYS.FOCUS
CONFIGURATION	CONF.CFG

* Only required when migrating system data.

** Only required if old repository is FOCUS.

- From the Data Repository Migrate drop-down menu, select Yes if the data repository for the previous release should be migrated. If Yes is selected, SMCNTRL, SMPRMTRS, and SMPRL data will be migrated. The default value is No.

If you select Yes, additional options become available.

The additional options let you migrate data in smaller portions, instead of all at once. The available options are:

Start Date

The starting date of the data to be migrated. The default value is the earliest date in the old repository.

End Date

The ending date of the data to be migrated. The default value is the latest date in the old repository.

- Accept the default dates, or enter a specific date that you want the migration to run.
- Click *Submit Job*. A confirmation of the request appears, stating that the deferred execution request is in the queue.
- On the ribbon, in the Repository group, click *Migration* and then click *View job*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and click *View job*.

The Deferred List pane opens, displaying the list of deferred requests and showing the request name that was displayed when the deferred migration job was scheduled. Examples of request names are listed in the *Defer ID* column.

The Status column shows the state of the deferred request, which can be Queued, Executing, or Ready. To see if the status has changed to Ready, either keep refreshing the page using the refresh options located at the top of the page, or redisplay the page at a later time. Once the status is Ready, right-click the request line and click Get to review the job output.

11. You can choose whether to remove the old configuration files or perform this action in the future.

On the ribbon, in the Repository group, click *Migration* and then click *Delete work files*.

or

From the navigation pane, right-click *Resource Management*, point to *Migration*, and click *Delete work files*.

- a. Click *Delete* and then *OK* to confirm.

Note: Once you choose to remove the old configuration files, the option to migrate data will no longer be available.

Event Routing

Event routing allows you to launch procedures or send email based on different server events. This feature includes five new pre-defined notifications having to do with five different types of events that can occur in Resource Governor. The notifications will be captured and visible on the edaprint and can be sent through email.

Error messages:

[36029 - Resource Management data repository validation failure](#)

Warning messages:

[36048 - Resource Governor canceled a request](#)

For more information on event routing and how to customize it, see the Event Routing section in the *Server Administration* manual.

Chapter 3

Rules Files and Knowledge Bases

Rules can be created in Resource Governor to set a variety of limits on database queries, such as establishing WHERE clause requirements or disallowing queries during specified time periods.

Resource Governor uses a language called Business Rule Language (BRL). BRL allows IF/THEN testing on certain information available when rules execute. Each BRL rule consists of three parts:

- The rule name, which is required, serves as a comment or description.
- A supporting condition (antecedent) or procedure statement. This is not syntactically necessary.
- A conclusion.

For more information about BRL, see [Business Rule Language \(BRL\)](#) on page 193.

In this chapter:

- [Overview](#)
 - [Working With Business Rules](#)
 - [Working with Resource Thresholds](#)
 - [Working with Knowledge Bases](#)
 - [Maintaining Knowledge Bases](#)
-

Overview

Resource Governor makes use of Knowledge bases to control which requests are run and which are not. A Knowledge base is built from business rules, data rules or both.

- Business rules are user coded using a wizard and/or editor and use information from the current environment. They do not use monitored data.
- Data rules are generated by combining historic usage data (monitored data) and user defined thresholds.

There are two types of Knowledge bases, Global and Data Source.

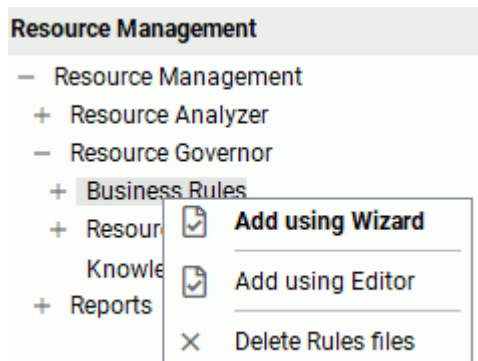
- ❑ A Global Knowledge base can only contain business rules, and there can only be one active at any given time. Its execution is dependent on the nature of the rules that it contains. If the rules are at the command level, for example disallow SELECT *, then execution is tied to which data request types are monitored, while a rule disallowing access to any data source that starts with the characters “ca” will always be executed.
- ❑ A Data Source Knowledge base is applied to individual data sources (synonyms or Direct Passthru requests) and therefore is only executed if the data source is referenced in a request. If multiple data sources are referenced in the request (JOIN or FROMs in an SQL statement), only one Knowledge base will be processed. It will be selected based on the first data source in the request list that is being actively governed.

Working With Business Rules

The left pane of the main Resource Management screen contains the Business Rules folder, which allows you to access the wizard. The wizard is a tool that allows you to create BRL statements accurately and quickly, using a dialog box that will generate rule syntax for you. The Business Rules folder also contains a plain text editing facility, where you can customize the rules or create a new rule rather than using the wizard.

Creating Business Rules

When you right-click the Business Rules folder, a context menu is available, as shown in the following image.



For additional information about building rules with BRL, see [Business Rule Language \(BRL\)](#) on page 193.

Procedure: How to Create a New Business Rule Using the Wizard

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. From the navigation pane, right-click *Business Rules*, and click *Add using Wizard*.

The Resource Governor New Business Rule pane opens, as shown in the following image.

Resource Governor New Business Rule

? Super User	No ▾
? Allow PRINT * or SELECT *	Yes ▾
? Number of WHERE or IF tests	No ▾
? Limit joined data sources	No ▾
? Cartesian product refinement	No ▾
? Disallow access to all data	No ▾
? Disallow access to a data source	No ▾
? Disallow Client type access	No ▾
Next ►	

The Resource Governor New Business Rule pane has the following options:

Super User

A user ID that is entered into the rules that will be allowed to execute any request that may be canceled by the governor under normal operating. A user ID may be up to 48 characters in length. Selecting Yes from the drop-down menu will open a text field for the user ID input.

Allow PRINT * or SELECT *

This setting allows the user to set whether or not a rule will be written. If Yes is selected, a rule will not be created and a SELECT * or PRINT * may be used at any time in a data request. If No is selected, a rule will be written disallowing the condition regardless of any other rules in the source. The default value is Yes.

Number of WHERE or IF tests

This setting allows you to specify the number, and name, of a condition.

- How many tests?** Allows you to specify that a certain number of conditions or relational clauses are required on a data request. Values can be 0-256. The default value is 0.
- With a field name?** Allows you to require a specific field name.

Limit joined data sources

The number of data sources used in the request. The limit is 16 and the default value is 0, or No. Selecting Yes will open a text field for input.

Cartesian product refinement

A Cartesian product occurs when two or more data sources are joined without any selection criteria. This produces an uncontrolled display of all possible connections between those tables. Requiring a WHERE or IF refinement prevents a Cartesian product from being issued. If Yes is selected, then one or more WHERE or IF clauses must be used to limit the data returned by the request that uses more than 1 data source. If No is selected, then all data for the request columns in the joined data sources will be returned. If you do not already have a rule that requires all requests to have a WHERE or IF test, you can select Yes. If an existing rule that requires all requests to have a WHERE or IF test is active, select No.

Disallow access to all data

Allows you to specify a time of day when data reporting is not allowed. The times are based on a 24-hour clock. The values allowed are 0000-2359.

Disallow access to a data source

Allows you to disable a data source name from reporting availability.

Disallow Client type access

The type of client being used to submit a request can be tested for, and if necessary the request can be canceled. For example, any request that is submitted from ReportCaster can be limited to specific times of the day or to a specific service of the server.

4. In this example:

- From the Allow PRINT * or SELECT * drop-down menu, select No.
- From the Limit joined data sources drop-down menu, Yes.

When Limit joined data sources is set to Yes, an additional box opens. Enter 3 in the box provided.

5. Click *Next*.

The New Business Rules window opens showing the generated code, as shown in the following image.

```

1  !
2  ! Rule Wizard Generated Template
3  !
4  RULE Select Asterisk
5  IF Select asterisk
6  THEN DBA Rules Concluded
7  AND Rule Number := 1
8  AND Run := "F"
9  AND Reason := "Select *"
10 AND Message1 := "No select asterisk."
11 !
12 RULE Allow Joins
13 IF Number of tables > 3
14 THEN DBA Rules Concluded
15 AND Rule Number := 2
16 AND Run := "F"
17 AND Reason := "Joins"
18 AND Message1 := "Only 3 joins allowed."
19 !
20

```

The New Business Rules window has the following buttons:

Verify Rules button 

Allows you to verify BRL syntax at any point while building rules. If there are no syntax errors, control is returned to the editor window with the rule text. If there are errors, a dialog displays error messages.

Note: The errors need to be corrected before you can continue.

Save As button 

Saves the new rule.

Select All

Selects all of the syntax in the window.

6. Click the Save As button.
7. Enter a name for the new business rule, (for example, *br1*), and click Save As.

Note: All business rule names in Resource Governor must be eight characters or less.

The new business rule is now listed in the expanded Business Rules folder in the navigation pane.

Procedure: How to Create a Business Rule Using the Editor

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. From the navigation pane, right-click *Business Rules* and click *Add using Editor*.

The New Business Rules window opens.

4. Enter the syntax for the new business rule and click the *Verify Rules* icon to check your code. Correct any errors and then click the *Save As* button.

For more information on building rules with BRL, see [Business Rule Language \(BRL\)](#) on page 193.

Procedure: How to Delete Rules Files

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. From the navigation pane, right-click *Business Rules* and click *Delete Rules files*.

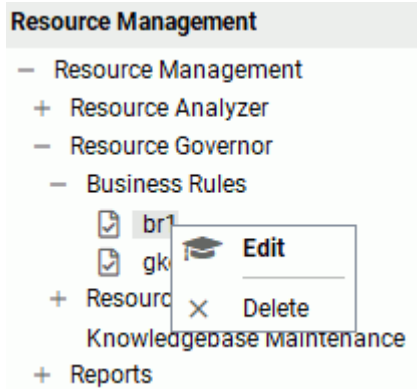
The Delete Business Rules files pane opens.

4. Select the rule you want to delete and click *Delete*.

The rule will be deleted from the Business Rules folder.

Editing Business Rules

When you right-click an existing rule, a context menu is available, as shown in the following image.



Procedure: How to Edit Existing Business Rules

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Business Rules* folder, right-click the rule you want to edit, and click *Edit*.

The Edit Business Rules pane opens, as shown in the following image.



```
1 !
2 ! Rule Wizard Generated Template
3 !
4 RULE Select Asterisk
5 IF Select asterisk
6 THEN DBA Rules Concluded
7 AND Rule Number := 1
8 AND Run := "F"
9 AND Reason := "Select *"
10 AND Message1 := "No select asterisk."
11 !
12 RULE Allow Joins
13 IF Number of tables > 3
14 THEN DBA Rules Concluded
15 AND Rule Number := 2
16 AND Run := "F"
17 AND Reason := "Joins"
18 AND Message1 := "Only 3 joins allowed."
19 !
20
```

4. Make any necessary changes, verify the rule changes by clicking the *Verify Rules* button, and then click the *Save* button.

Note:

- A rules file that is being used in one or more Knowledge bases cannot be changed. Any changes have to be saved to a different name, or the Knowledge base files have to be deleted first.
- If a rule file that is in use is edited, a warning window opens and the *Save* button is not available. Click *OK* to continue.
- To save the rule under a different name, click the *Save As* button and enter a new name.
- When a business rule is edited, it will always open in the *Edit Business Rules* window, regardless of how it was originally created.

Procedure: How to Delete an Existing Rule

1. Launch the Web Console.

2. Click the *Resource Management* option on the sidebar.
3. Expand the *Business Rules* folder, right-click the rule you want to delete, and click *Delete*.
A confirmation window opens.
4. Click *OK*.

The rule will be deleted from Resource Management.

Note: A rules file that is being used in one or more Knowledge bases cannot be deleted. The Knowledge base files have to be deleted first.

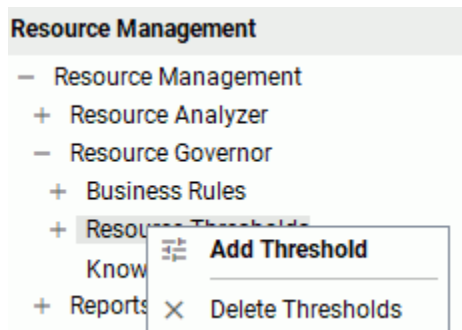
Working with Resource Thresholds

Threshold parameters are used to build data rules after monitored data retrieval. Threshold parameters enable you to create shifts, delineated by days, hours, and times, that control requests made by end users. They are pertinent to data rules created from historical monitored data because they contain the thresholds that are compared to the historically gathered resources used.

In addition, threshold parameters enable you to set thresholds. Thresholds are the maximum resources that are allowed for a shift. You can set thresholds for CPU seconds, number of rows returned in a report, elapsed time in seconds, and input/output units.

Creating Thresholds

When you right-click the Resource Thresholds folder, a context menu is available, as shown in the following image.



Procedure: How to Create Thresholds

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.

- From the navigation pane, right-click *Resource Thresholds* and click *Add Threshold*.

The Add Resource Threshold pane opens, as shown in the following image.

Add Resource Threshold

? Name

? Status

? Thresholds Elapsed Time (Seconds)
 CPU Time (Seconds)
 Output Rows
 I/O Used

? Time Range Start (format: HH:MM)
 End (format: HH:MM)

? Date Type Date Range
 Day of Week

? Date Range Start (format: MM/DD)
 End (format: MM/DD)

The Add Resource Threshold pane has the following fields and options:

Name

The name of the threshold.

Status

The state of the threshold indicates whether the parameter is making rules. If set to Active, the parameter will be included in any Knowledge base files that are built. The default value is Active, with the time and date ranges specified. If set to Inactive, the parameter will not be included in the built rule file.

Thresholds

The following threshold values are available to be set:

- Elapsed Time.** The maximum number of wall clock seconds that a request can use without being canceled.
- CPU Time.** The maximum number of CPU seconds that a request can use without being canceled.
- Output Rows.** The maximum number of rows that a request can potentially return without being canceled.
- I/O Used.** The maximum number of input/output units a request can generate without being canceled.

Time Range

The times when this set of thresholds will be in effect.

Date Type

Indicates what days of the week the threshold will be in effect for. Select the *Day of Week* radio button to set specific days of the week. To set a range, select the *Date Range* radio button and enter a Start and End date in the MM/DD format. The dates do not use a year since the whole parameter may be used across years without modification needed.

4. Click *Add*.

Procedure: How to Delete Thresholds

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Right-click the *Resource Thresholds* folder and click *Delete Thresholds*.

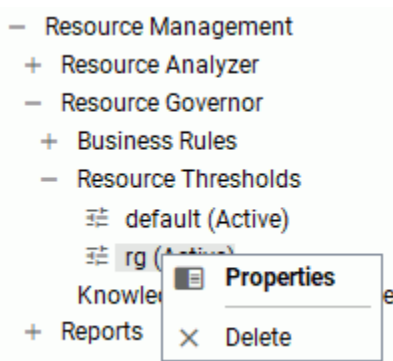
The Delete unused Resource Thresholds pane opens.

4. Select the thresholds you want to delete and click *Delete*.

The Delete Resource Thresholds - Status pane opens, showing which thresholds were deleted.

Editing Resource Thresholds

When you right-click an existing threshold, a context menu is available, as shown in the following image.



Procedure: How to Edit a Threshold

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Resource Thresholds* folder.
4. Double-click the threshold you wish to edit, or right-click it and click *Properties*.

The Resource Threshold Properties pane opens.

5. Update the properties as necessary.

Note: Time Range and Date Range cannot be updated in the Resource Threshold Properties pane. In order to change these settings, you must make a new threshold. The only fields that can be changed from the this properties pane are Status, Thresholds, and Day of Week.

6. Click *Update*.

Procedure: How to Disable Thresholds

If you select a value of Inactive for the Status of a shift, rules will not be activated for it.

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Resource Thresholds* folder.
4. Right-click a threshold and click *Properties*.

The Resource Threshold Properties pane opens.

5. Select *Inactive* from the Status drop-down menu.
6. Click *Update*.

Procedure: How to Delete a Threshold

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Resource Thresholds* folder.
4. Right-click the threshold you wish to delete and click *Delete*.

A confirmation dialog box opens.

5. Click *OK*.

The threshold has been removed from the Resource Thresholds folder.

Working with Knowledge Bases

Knowledge bases can be applied globally, in which case only one can be active at a time, and/or applied to an individual data source (synonym or Direct Passthru source).

Applying a Global Knowledge Base

Note: You must create a business rule file before adding a global Knowledge base. For more information on creating rules, see [Creating Business Rules](#) on page 104.

Procedure: How to Add a Global Knowledge Base

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, in the Configuration group, click *Configuration*, point to *Global Governing*, and then click *Add Knowledgebase*.

The Add Global Knowledgebase pane opens, as shown in the following image.

Add Global Knowledgebase

The image shows a dialog box titled "Add Global Knowledgebase". It contains a "Business Rules" label with a question mark icon and a dropdown menu showing "- Select -". Below this is a "Knowledgebase" label and a text input field. At the bottom, there are two buttons: "Add" and "Cancel".

4. Select the rule file you want to use from the Business Rules drop-down menu.
5. Enter a name for the Knowledge base.
Note: The Knowledge base name cannot contain a space or special characters.
6. Click *Add*.

The List Global Knowledgebase pane opens, listing the new Knowledge base. The default status of the new global Knowledge base is Off. You can change it by right-clicking the status and selecting an option from the context menu.

The new Knowledge base now appears in the navigation pane, under the Knowledgebase Maintenance folder.

Procedure: How to List Global Knowledge Bases

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, in the Configuration group, click *Configuration*, point to *Global Governing*, and then click *List Knowledgebases*.

The List Global Knowledgebase pane opens, as shown in the following image.

List Global Knowledgebase

Close

Note: Knowledgebase names that are red cannot be found and must be rebuilt. Right click on the name and select Rebuild.

Knowledgebase Name	Build Status	Date Created	Govern Status	Rules File Name
RG	0	03/28/2018	Off	BR1
RGBR1	0	03/28/2018	Off	GKECR

The columns shown have the following context menu options available:

Knowledgebase Name Column

- View.** Opens the View Global Business Rules pane that enables you to view the business rule associated with that Knowledge base.
- Rebuild.** Rebuilds the selected Knowledge base.

Note: If the Knowledgebase Name is shown in red, it indicates that the Knowledge base was not found and must be rebuilt using the *Rebuild* option.

Govern Status Column

- On.** Turns governing on for the selected Knowledge base. The request will be canceled if the Knowledge base rules are applicable to the request.
- Off.** Turns governing off for the selected Knowledge base.
- Advise.** Activates governing for the selected Knowledge base, but allows the request to execute after a message is issued indicating the request would have been canceled if governing was set to On.

Procedure: How to Activate Advise Message Display in WebFOCUS

To activate the Advise message display, you must go to the WebFOCUS Administration Console.

1. Go to the URL:

`http[s]://machine:port/context/admin`

where:

machine

Is the network ID of your computer.

port

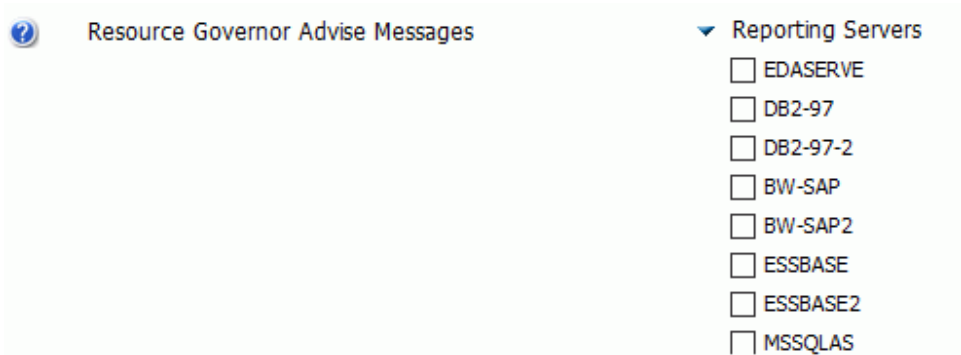
Is the number of the port that connects your computer to the server hosting WebFOCUS.

context

Is the local address for WebFOCUS. For example, *ibi_apps*.

2. On the Sign In page, type the *User Name* and *Password* of a user that has privileges to open the Administration Console, and click *Sign In*.
The Administration Console opens automatically.
3. In the Application Settings folder, click *Client Settings*.
The Client Settings open in the main configuration page.
4. Find the Resource Governor Advise Messages setting and expand the *Reporting Servers* option.

A list of available servers opens, as shown in the following image.



5. To enable the display of Advise messages for a specific server, select the checkbox for the server.
6. Click Save.

Procedure: How to Delete Global Knowledge Bases

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. On the ribbon, in the Configuration group, click *Configuration*, point to *Global Governing*, and then click *Delete Knowledgebases*.

The Delete Global Knowledgebases pane opens.

4. Select the Knowledge bases you want to delete and click *Delete*.

Applying an Individual Knowledge Base

Note: Before applying an individual Knowledge base, you must monitor a data source. For more information, see [How to Monitor Individual Data Sources](#) on page 75.

Procedure: How to Add an Individual Knowledge Base

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. From the navigation pane, right-click the monitored data source, point to *Governing*, and then click *Add Knowledgebase*.

The Add Knowledgebase pane opens, as shown in the following image.

Add Knowledgebase for ibisamp/car

? Knowledgebase

? Type

? Earliest Monitored Date:

Latest Monitored Date:

? Time

Note: If the object you selected has no monitored data, the Add Knowledgebase pane will show fewer fields, as shown in the following image.

Add Knowledgebase for ibisamp/brokers

? Knowledgebase

? Type

? Business Rules

The following options are available:

Knowledgebase

The name of the new Knowledge base.

Note: The Knowledge base name cannot contain a space or special characters.

Type

The rule type. The following options are available:

- Data.** Available for a data source that has had historical monitoring at some time.
- Business.** Only allows a business rule source to be included in a built Knowledge base file. A list of those rule sources will be available.

- Both.** Allows data and business rules to be included in the Knowledge base file.

Earliest Monitored Date

Shows the earliest monitored date for this data source.

Latest Monitored Data

Shows the latest monitored date for this data source.

Time

A rule type of Data or Both allows for either CPU or Elapsed times to be used as benchmarks in the monitored data as compared to thresholds entered through governing parameters or shifts. Elapsed time is wall clock time for a data request to run as opposed to the CPU time used in processing.

Business Rules

A list of the available business rules. This drop-down menu is only available if the object you selected has no data.

- Click *Add*.

The List Knowledgebase pane opens, listing the new Knowledge base.

The new Knowledge base now appears in the navigation pane, under the Knowledgebase Maintenance folder.

Procedure: How to List Individual Knowledge Bases

- Launch the Web Console.
- Click the *Resource Management* option on the sidebar.
- From the navigation pane, right-click a monitored data source, point to *Governing*, and then click *List Knowledgebases*.

The Knowledgebase pane opens, as shown in the following image.

Knowledgebases for ibisamp/car

Close

Note: Knowledgebases that are red cannot be found and must be rebuilt. Right click on the name and select Rebuild.

Knowledgebase Name	Build Status	Date Created	Govern Status	Rules Type	Rules File Name	Time Type	Monitored Data Dates
CAR RULE1	0	03/28/2018	Off	Data		CPU	03/28/2018-03/28/2018
CAR RULE2	0	03/28/2018	On	Both	GKECR	CPU	03/28/2018 - 03/28/2018

The columns shown have the following context menu options available:

Knowledgebase Name Column

- View.** Opens the View Knowledgebase Rules pane that enables you to view the business rule associated with that Knowledge base.
- Rebuild.** Rebuilds the selected Knowledge base.

Govern Status Column

- On.** Turns governing on for the selected Knowledge base. The request will be canceled if the Knowledge base rules are applicable to the request.
- Off.** Turns governing off for the selected Knowledge base.
- Advise.** Activates governing for the selected Knowledge base but allows the request to execute after a message is issued indicating that the request would have been canceled if governing was set to On. For more information, see [How to Activate Advise Message Display in WebFOCUS](#) on page 117.

***Procedure:* How to Delete Individual Knowledge Bases**

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. From the navigation pane, right-click a monitored data source, point to *Governing*, and then click *Delete Knowledgebases*.

The Delete Knowledgebase pane opens.

4. Select the Knowledge bases you want to delete, and click *Delete*.

Maintaining Knowledge Bases


When you right-click an existing Knowledge base, a context menu is available, as shown in the following image.



Procedure: How to View a Knowledge Base Rule

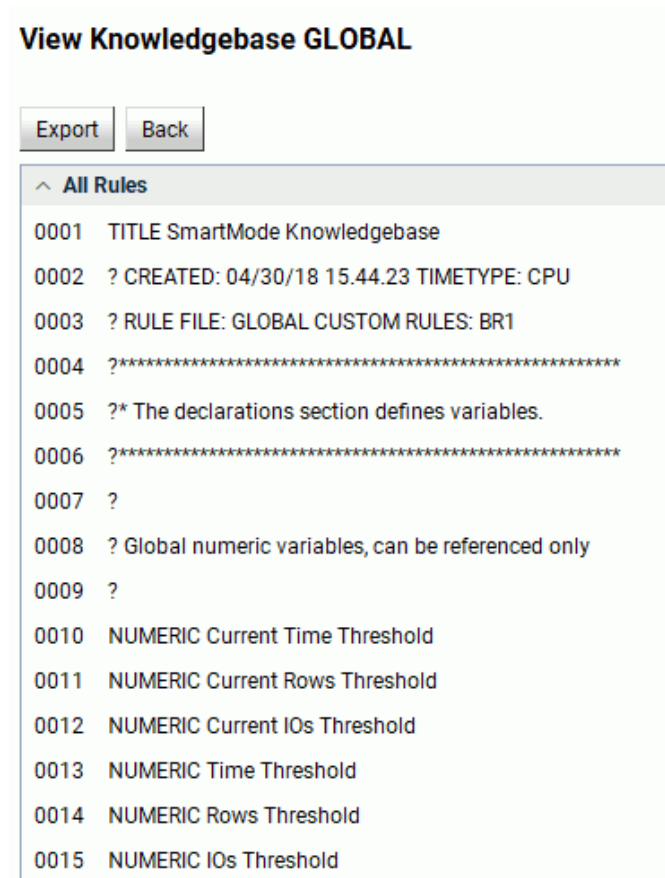
1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Knowledgebase Maintenance* folder.
4. Right-click the Knowledge base and click *View Rules*.

The View Knowledgebase pane opens, as shown in the following image.



5. Click *View All Rules* to export the selected Knowledge base.

The View Knowledgebase pane opens showing all business rules, as shown in the following image.



6. Click *Export*.

A dialog box opens that indicates the file location of the exported file.

7. Click *OK* to close the dialog box.

Procedure: How to Rebuild the Knowledge Base

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Knowledgebase Maintenance* folder.
4. Right-click the Knowledge base and click *Rebuild*.

The Rebuild Knowledgebase pane opens, as shown in the following image.

The screenshot shows a window titled "Rebuild Knowledgebase CAR_RULE2 for ibisamp/car". It contains several input fields and buttons:

- Knowledgebase:** CAR_RULE2
- Type:** Data (dropdown menu)
- Earliest Monitored Date:** 03/28/2018 (calendar icon)
- Latest Monitored Date:** 03/28/2018 (calendar icon)
- Time:** CPU (dropdown menu)
- Buttons:** Rebuild, Cancel

Note: If the object you selected has a different rule type, the Rebuild Knowledgebase pane will show fewer fields.

The Rebuild Knowledgebase pane has the following options:

Knowledgebase

The name of the selected Knowledge base.

Type

The rule type. The following options are available:

- Data.** Available for a data source that has had historical monitoring at some time.
- Business.** Only allows a business rule source to be included in a built Knowledge base file. A list of those rule sources will be available.
- Both.** Allows data and business rules to be included in the Knowledge base file.

Time

The type of time resource used in the data rule. This option is only available if you select *Business* or *Both* from the Type drop-down menu.

Business Rules

Shows all of the business rule sources that are available. This option is only available if you select *Business* or *Both* from the Type drop-down menu.

5. Click *Rebuild*.

Procedure: How to Delete a Knowledge Base

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Knowledgebase Maintenance* folder.
4. Right-click the Knowledge bases you want to delete and click *Delete*.
5. Click *OK* in the confirmation window.
6. To delete more than one Knowledge Base at a time, right-click *Knowledgebase Maintenance* and click *Delete Knowledgebase files*.

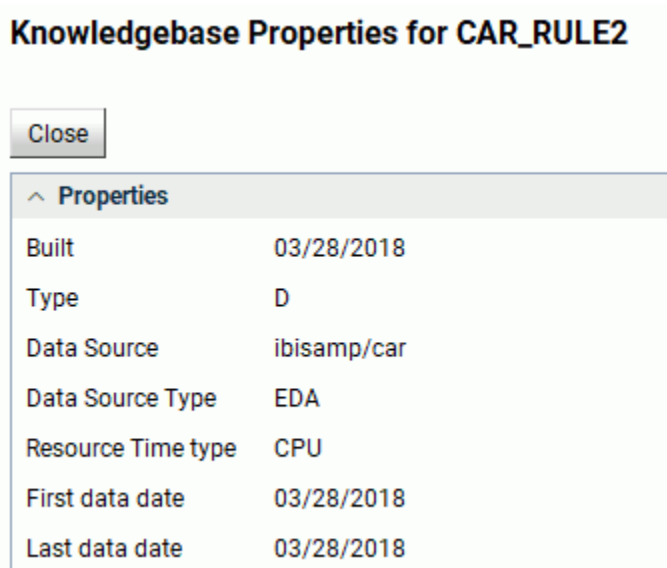
The Delete Knowledgebase files pane opens.

7. Select the Knowledge bases you want to delete and click *Delete*.

Procedure: How to View Knowledge Base Properties

1. Launch the Web Console.
2. Click the *Resource Management* option on the sidebar.
3. Expand the *Knowledgebase Maintenance* folder.
4. Right-click a Knowledge base and click *Properties*.

The Knowledgebase Properties pane opens, as shown in the following image.



The Knowledgebase Properties pane has the following fields:

Built

The date the Knowledge base was built.

Type

The rule type.

Business Rules

The business rule source being used by the Knowledge base.

Data Source

The data source used for the Knowledge base.

Data Source Type

The type of data source.

Resource Time type

The type of time used as benchmarks in the monitored data.

First Data Date

The first date of monitored data for the data source.

Last Data Date

The last date of monitored data for the data source.

5. Click *Close* to exit the Knowledgebase Properties pane.

Resource Governor Report Options

After you have enabled Usage Monitoring for a period of time and Resource Governor has collected a reasonable amount of data, you can display several reports and graphs about the monitored data. This chapter describes the various reports you can run and view with Resource Governor.

In this chapter:

- [Reports Overview](#)
 - [Resource Governor Reporting](#)
 - [Resource Governor Report Options](#)
 - [General Reports](#)
 - [Resource Governor Reports](#)
-

Reports Overview

Resource Governor produces a range of different reports, depending on the type of information you want and the types of issues you are addressing:

- Monitored Sessions.** Shows the number of user sessions by date.
- Monitored Commands.** Shows the number of connections and the amount of resources used.
- Repository Statistics.** Shows the total number of records collected during the time period specified.
- Data Sources Never Used.** Opens the Unused Master File Descriptions report, which lists the data sources which have never been used.
- Procedures Never Used.** Opens the Procedures Never Used report, which lists the procedures in the APP PATH which have never been executed.
- ACI Monitoring.** Opens the ACI Summary report, which provides information about changes to the server configuration files, and user application files. Actual changes made are not reported, but who, when, and the type of access is shown.

The Resource Governor Reports folder contains the following reports:

- ❑ **Monitored Data Source.** Provides an overview of the data sources that are being monitored, including number of requests, and the first and last access dates.
- ❑ **Monitored Procedures.** Provides an overview of the procedures that are being monitored, including the number of executions, and the first and last execution dates.
- ❑ **Resource Thresholds.** Shows parameter details of all Resource Governor thresholds that have been built.
- ❑ **Knowledgebase Properties.** Shows details of all the knowledge bases that have been built.
- ❑ **Threshold Advisor.** Recommends thresholds based on a statistical model.
- ❑ **Governed Requests.** Shows which queries have been governed or canceled and the rule file used to govern.

Resource Governor Reporting


For all platforms, Resource Governor reporting is accessed from the Web Console.

Procedure: How to Access Resource Governor Reports

To access Resource Governor reports:

1. Launch the server Web Console.

2. Click *Resource Management* on the sidebar and expand the *Reports* and folders, as shown in the following image.



Procedure: How to Run Resource Governor Reports

To run Resource Governor reports:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree and click *Run*.
3. The Report Filter window opens. Make any desired changes and click *View Report*. For more information, see [Report Filters](#) on page 138.

Procedure: How to Submit a Report

To submit a report:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree and click *Submit*.

The Report filter for Repository Statistics pane opens.

3. Make any desired changes from the filtering options provided. For more information on these options, see [Report Filters](#) on page 138.

4. From the calendar icon and the drop-down menu, specify a date and time to execute the report.
5. Click *Submit Report* to submit the report to the deferred queue.

Procedure: How to View a List of Submitted Reports

To view a list of submitted reports that have been run in deferred mode:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree and click *View Submitted*.

The Deferred List pane opens, showing a list of submitted reports.

If you right-click an item in the list, the following options are available:


- View.** Opens the report.
- Download.** Saves the selected report to your Downloads folder. In Windows, the default location of this folder is C:\users*username*\Downloads.
- Delete.** Deletes the selected item.

Procedure: How to Schedule a Report

You can schedule reports to run and be automatically distributed. To schedule a report:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree, point to *Schedule and E-Mail*, and click *Add job*.

The Report filter for the selected report opens, as shown in the following image.



3. Select your desired filtering conditions and click *Next*.

The Manage Schedule and E-Mail pane opens, as shown in the following image.

The screenshot shows a configuration pane with the following sections and options:

- General**:
 - ? Schedule Status: **No Schedule** (dropdown) Enables/Disables Schedule
- Log and output destinations**: (collapsed)
- E-Mail On Start**: (collapsed)
- E-Mail On Completion**:
 - ? Enable:
 - ? Mail to:
 - To**: [dropdown]
 - Cc**: [dropdown]
 - Bcc**: [dropdown]
 - Use a semicolon as the separator character for multiple email addresses
 - ? Advanced:
- E-Mail On Failure**: (collapsed)

Buttons: **Cancel** and **Set**

The Manage Schedule and E-Mail pane has the following options:

Schedule Status

Determines if there is a schedule and if so, whether the schedule is active or inactive.

Log and output destinations

These options control where the log information from the job should be sent when the job finishes.

The following options are available:

- ETLLOG**. The log is sent to the DataMigrator log.
- EDAPRINT**. The log is sent to the Server log.
- E-Mail**. The log is sent as an attachment to an email message if specified in Job Properties.

- Send HTML output as inline message.** When selected, if any stored procedures called by the flow generate an output, the message is included as message text instead of as an attachment. This option is only available in the E-Mail On Completion section.

Multiple options can be selected. ETLLOG and E-Mail are selected by default.

Note:

- If the job is run from the DMC or Web Console, information will go to the console log regardless of this setting.
- The *E-Mail On Start* and *E-Mail On Completion* options are not relevant to the Resource Management Archive job and should not be used.

Enable

Select this to enable when an email should be sent.

Mail to

The email addresses where the messages will be sent. Place each email address (or user ID configured on the server with an associated email address) on a new line with no separators.

Advanced

Select this option to get the following options:

- Mail to addresses defined in procedure.** You can send the messages to email addresses defined in the procedure by either typing in the location of the procedure, or clicking the ellipsis button and manually selecting it from the Select Procedure dialog box.
- Importance.** Select the importance for the e-mail from the drop-down menu. Choices are Low, Normal, or High.
- Subject.** The subject line of the message.
- E-Mail Message.** The body of the message.

4. If you select *Active* or *Inactive* from the Schedule Status drop-down menu, the Manage Schedule and E-Mail for procedure pane opens with more fields, as shown in the following image.

^ General	
? Schedule Status	Active <input type="text"/> Enables/Disables Schedule
? Schedule Type	Run Once <input type="text"/> Determines whether a flow will execute
^ Start and stop range	
? Start Date	05/07/2018 <input type="text"/> Specifies the start time and date for the execution of
? Start Time	12 <input type="text"/> 27 <input type="text"/> Specifies the start time and date for the execution of
^ Special date/time ranges	
? Specify Stop Time	<input type="checkbox"/> Specifies forced end time
^ Log and output destinations	
^ E-Mail On Start	
^ E-Mail On Completion	
? Enable	<input type="checkbox"/>
? Mail to	To <input type="text"/>
	Cc <input type="text"/>
	Bcc <input type="text"/>
	Use a semicolon as the separator character for multiple email addresses (ic
? Advanced	<input type="checkbox"/>
^ E-Mail On Failure	
<input type="button" value="Cancel"/>	<input type="button" value="Set"/>

The Manage Schedule and E-Mail for procedure pane contains the following additional fields and options:

Schedule Type

Determines whether a process flow will execute once, on a recurring basis, on several specific days (multi-day), or when the server starts.

Interval Type

Specifies the interval in which a process flow will be executed. For example, if you select Month as your Interval Type, and select 2 as your Interval Number, the flow will execute every two months. This option is only available if the schedule type is set to *Recurring*.

Interval Value

Specifies the frequency at which a process flow will be executed. For example, if you select Week as your Interval Type, and select 2 as your Interval Value, the flow will execute every two weeks. This option is only available if the schedule type is set to *Recurring*.

Start Date

Specifies the start date for the execution of the process flow.

Start Time

Specifies the start time for the execution of the process flow using 24-hour time.

Specify Stop Time

Allows you to specify an end time for the execution of the process flow using 24-hour time. This option is only available if the schedule type is set to *Run Once* or *Multi-Day*.

Specify End Date

Allows you to specify an end date and time for the execution of the process flow. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

End Date

Specifies the end date for the execution of the process flow. This option is available if Specify End Date is selected.

End Time

Specifies the end time for the execution of the process flow using 24-hour time. This option is available if Specify End Date is selected.

Days of the Week

Specifies the day of the week on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Days of the Month

Specifies the day of the month on which the process flow will be executed. This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

Intraday Start

Specifies the beginning of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Intraday End

Specifies the end of the daily time span of the process flow using 24-hour time. This option is only available if the schedule type is set to *Recurring* and the interval type is set to time, such as *Minutes*.

Run on Holidays

Specifies the option to execute the process flow on desired holidays.

The following options are available:

- Skip.** Excludes run on desired holidays.
- Only.** Only run on desired holidays.

This option is only available if the schedule type is set to *Recurring* or *Multi-Day*.

5. Click *Set*.

A confirmation dialog box opens.

6. Click *OK*.

Procedure: How to View Scheduled Reports

To view the details of a scheduled report:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree, point to *Schedule and E-Mail*, and click *List job*.

The Scheduled jobs pane for the selected report opens, as shown in the following image.

Scheduled jobs - Monitored Commands

Job Name	Schedule	Schedule type
rmsr0004	Active	Run Once.

If you right-click a drill-down item in a column, the following options are available:

Job name Column

- Filters.** Opens the Report filter pane for that enables you to update the filters for the selected report.
- View Reports.** Opens a pane that shows when the scheduled report was run. If you right-click an item in the list, the following options are available:
 - View.** Opens the report.
 - Download.** Saves the selected report to your Downloads folder. In Windows, the default location of this folder is C:\users*username*\Downloads.
 - Delete.** Deletes the selected item.

Schedule Column

- Manage Schedule.** Opens the Manage Schedule and E-Mail for procedure pane.

Procedure: How to Delete Scheduled Reports

To delete a report that has been scheduled to run:

1. Expand the *Reports* folder on the tree in the navigation pane.
2. Right-click the desired report in the tree, point to *Schedule and E-Mail*, and click *Delete job*.
The Delete Scheduled Reports pane for the selected report opens.
3. Select the scheduled jobs you want to delete by selecting the checkbox next to its name. To select all listed reports, select the checkbox in the header.
4. Click *Delete*.

Resource Governor Report Options

There is general information that applies to all Resource Governor reports. This information includes setting date ranges for report selection criteria, drilling down on reports for more details, displaying reports in graphical format, and the online help.

Report Filters

Before you execute a report, you can select the date range of usage monitoring data to include along with other criteria used in the report. To specify these options, right-click the report you want to run and click *Run*.

The Report Filter window opens, as shown in the following image.

Report filter for Monitored Commands

Select server: FUENT-PC:8121

Report format: HTML

Visualize Data: Yes

Report period: Date range

Start Date: 2019/2/19

End Date: 2019/2/22

View Report

The above image is an example of the filter pane. Different reports may have different options available. See *Performance Analysis Reports*.

All other reports will only have the first three options available.

The following options are available:

Select server

The server, or servers, you want to report on. Multiple server support is only available if the server you are connected to is part of a many-to-one configured repository.

Visualize Data

Allows you to add visualization to numeric resource values, such as CPU and duration times. This option is enabled by default.

Report format

Allows you to select the output format of the report. The available options are HTML, active report, PDF, Excel, and PowerPoint. The default value is HTML.

Report period

The time range that the report will be run on.

Start Date

Allows you to select the first date of the reporting range.

End Date

Allows you to select the first date of the reporting range. The default end date is the current date.

The information you supply here will apply to any report you run. To change this information, return to this screen by running another report and changing the filters.

Resource Governor report selection criteria information is saved for the same browser session, so you will not need to reset it when you reconnect to the server.

Procedure: How to Turn Off Data Visualization

You have the option of turning off data visualization when displaying reports. By default, the Visualize Data checkbox is selected, indicating that data visualization is on.

To turn off data visualization:

1. Clear the *Visualize Data* checkbox.
2. Click *View Report*.

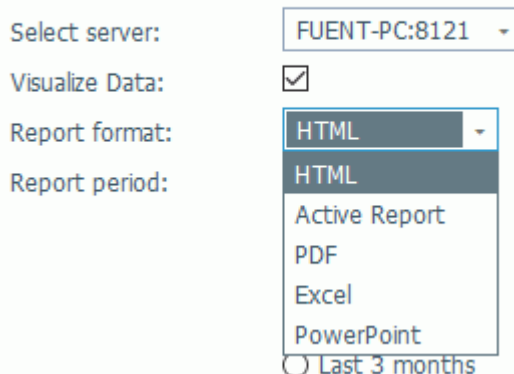
Procedure: How to Format a Report

To change the report format:

1. Click the *Report format* drop-down menu.

The following image shows all of the report options available. If you are licensed for active reports, you can choose that as your report format. For more information on active reports, see the *WebFOCUS Managed Reporting End User's Manual*.

Resource Management report filter



The screenshot shows the 'Resource Management report filter' interface. It includes the following elements:

- Select server:** A text box containing 'FUENT-PC:8121' with a dropdown arrow.
- Visualize Data:** A checkbox that is checked.
- Report format:** A dropdown menu that is open, showing the following options: HTML (highlighted), HTML, Active Report, PDF, Excel, and PowerPoint.
- Report period:** A radio button next to the text 'Last 3 months'.

2. Select a report format from the drop-down menu. The following options are available:

- HTML
- active report
- PDF
- Excel
- PowerPoint

3. Click *View Report*.

Note: Both Excel and PowerPoint reformatted reports are accessed using the download link in your browser. Drill down links in PowerPoint reports will only be active if PowerPoint is in display mode. Additional pop-up messages may be displayed by these two products when following the drill-down links. These messages are outside of the control of the Resource Management product.

***Procedure:* How to Set Selection Criteria Dates for Reports**

By default, suggested dates are shown in the Start and End Date input boxes. You may change these values by typing directly into the boxes, or by selecting dates from a pop-up calendar. To set selection criteria dates for reports:

1. Type a date into the date input box on the left (Start Date).

or

Click the that appears to the right of the date input box.

A calendar appears from which you can choose the month, day, and year for the Start Date of the data on which you want to report.

2. Type a date into the End Date input box or select a date from the pop-up calendar. If no date is selected, the current date will be used by default.

3. Click *View Report*.

Online Help

When you run any of the reports, you have an online help feature at your disposal, accessible through the ? *Help* button in the report window. The online help provides background information about the report you are running.

For more information about the help features, see [Getting Help](#) on page 153.

Drilling Down With Reports

Initially, when you execute reports, you will usually see a general summary report for the category. Most of the column headings of the reports are hyperlinked. Click the column heading links to re-sort the report by different columns. In addition, other hyperlinks in the report let you drill down to more detailed information. These hyperlinked drill-down options are described in more detail throughout this chapter.

The following sections describe the individual reports. Some of these sections contain a table of drill-down links. In these tables, terms shown in *italic* represent placeholders, indicating the term can have different values, depending on the selections for that report.

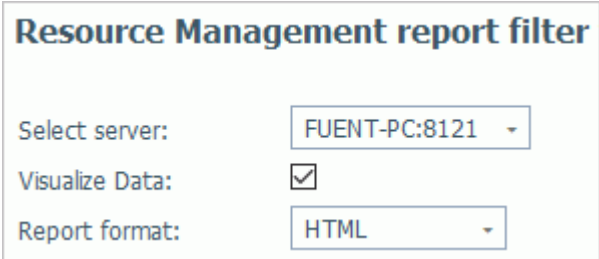
Monitor Preference and Reports

If you change the Monitor Preference setting and do not collect all the information, some reports will have missing data. When those reports are executed, a message will be displayed informing the user that no data is available.

Alternate Reporting Output

Many reports can be output in both tabular and graphical formats. Reports can also be reformatted into a number of different outputs.


Reports initially appear in the format specified by the Report format setting in the Resource Management window, as shown in the following image.



The image shows a dialog box titled "Resource Management report filter". It contains three settings:

- Select server:** A dropdown menu with "FUENT-PC:8121" selected.
- Visualize Data:** A checked checkbox.
- Report format:** A dropdown menu with "HTML" selected.

Once in the report, you can change its format by selecting a different option from the drop-down menu located on the upper left of the report page, as shown in the following image.



The drop-down menu lists the different report outputs, allowing you to change the display of the current report. Reports run from the tree will use the Report format set in the Resource Management: Report Filter window.

General Reports

These reports provide overviews of the data sources and procedures being monitored, as well as those procedures in the server application path that have never been executed. General Reports appear only on the Web Console.

To access any of the reports in the Web Console, click *Resource Management* on the sidebar and expand the *Reports* folder. Right-click the desired report, and click *Run* from the context menu. After specifying report filters, click *View Report*.

Monitored Sessions

The Monitored Sessions report provides an overview of the procedures and commands that are being monitored, including the number of records processed, and the rows returned.

The following image shows the Monitored Sessions report.

Session Summary Report
Server Name: FUENT-PC:8121
Date Range: 05/12/2015 - 09/12/2016

Date	Session Count	Duration	CPU Time	Wait Time	Adapter Time	Monitored Procedures	Monitored Commands	I/O Used	Records Processed	Rows Returned	Transactions
08/12/2016	3	00:00:00.3250	00:00:00.2800	00:00:00.0450	00:00:00.0310	0	9	10	1,058	12	0
08/31/2016	32	00:00:24.6340	00:00:01.2010	00:00:23.4330	00:00:08.7320	0	10	0	297	297	228
09/09/2016	55	00:00:14.6630	00:00:02.9570	00:00:11.7060	00:00:03.5900	8	27	40	3,249	906	299
09/12/2016	32	00:00:19.9240	00:00:01.6400	00:00:18.2840	00:00:11.4340	0	13	10	1,405	359	296
Grand Total	122	00:00:59.5460	00:00:06.0780	00:00:53.4680	00:00:23.7870	8	59	60	6,009	1,574	823

This report has one hyperlink that allows you to drill down to other reports, as described in the following table.

Click a hyperlink in the:	Report Generated	Description
Date column	Session Summary Report by Hour	Lists the procedures and commands that are being monitored by hour. Allows you to drill down to see this report by quarter, and then by detail.

Monitored Commands

The Monitored Commands report provides an overview of the resources being used by each command, including the execution, CPU, and wait time.

The following image shows the Monitored Commands report.

Resources Used by Command
 Server Name: [FUENT-PC:8121](#)
 Date Range: 08/12/2015 - 09/12/2016

Date	Command	Usage Count	Execution	CPU Time	Wait Time	Adapter Time	I/O Used	Records Processed	Rows Returned	Transactions
09/12/2016	TABLE	7	00:00:00.3240	00:00:00.0460	00:00:00.2780	00:00:00.0000	0	173	173	0
	TABLEF	7	00:00:04.2640	00:00:00.2660	00:00:03.9980	00:00:00.0000	0	206	206	0
09/09/2016	JOIN	1	00:00:00.0000	00:00:00.0000	00:00:00.0000	00:00:00.0000	0	0	0	0
	TABLE	10	00:00:00.7900	00:00:00.1120	00:00:00.6780	00:00:00.1230	2	201	201	0
	TABLEF	9	00:00:02.6030	00:00:00.1800	00:00:02.4230	00:00:00.0000	0	208	208	0
08/31/2016	TABLE	7	00:00:00.5270	00:00:00.0460	00:00:00.4810	00:00:00.0000	0	133	133	0
	TABLEF	3	00:00:09.9270	00:00:00.3270	00:00:09.6000	00:00:00.0000	0	162	162	0
08/12/2016	TABLE	8	00:00:00.1030	00:00:00.0550	00:00:00.0480	00:00:00.0300	8	1,058	12	0
	TABLEF	1	00:00:00.0050	00:00:00.0000	00:00:00.0050	00:00:00.0000	2	0	0	0
Grand Total:			00:00:18.5430	00:00:01.0320	00:00:17.5110	00:00:00.1530	12	2,141	1,095	0

Percentage and Ratio fields are rounded values
 Resource Analyzer
 1

Date: 09/12/16 Time: 15.16.37

Repository Statistics

The Repository Statistics report provides an overview of the total number of records collected during the time period specified.

The following image shows the Repository Statistics report.

RMLDB Repository record counts															
By Year and Month															
Server Name: FUENT-PC:8121															
Date Range: 02/01/2018 - 02/22/2018															
Server Name	Year	Month	Sessions	Procedures	Procedure Lines	WebFOCUS	Query	Query Lines	Governor	Remotes	Data Sources	Fields	Functions	Relations	Audits
FUENT-PC:8121	2018	February	337	26	162	0	802	88	0	0	0	0	0	0	0

Data Sources Never Used

The Data Sources Never Used report provides an overview of unused data sources, including the application directory they are found in, and the last modification date and time.

The following image shows the Data Sources Never Used report.

Unused Master File Descriptions				
Server Name: FUENT-PC:8121				
Application	Data Source	Last Modification Date	Last Modification Time	
baseapp	dmortmplx	2012/05/23	12.41.43	
	dmpers	2010/09/28	14.23.08	
	dmxduplg_dupl	2012/05/22	15.40.14	
	dmxord	2012/05/23	11.29.42	
dimensions	ed_type_lookup	2016/05/23	15.45.56	
	ed_type_lookup	2016/06/28	15.57.55	
	wf_retail_age	2016/05/23	15.45.06	
	wf_retail_age	2016/06/28	15.57.49	

Procedures Never Used

The Procedures Never Used report lists those procedures in the application path of the server that have never been executed.

The following image shows the Procedures Never Used report.

Unused Procedures		Server Name: FUENT-PC:8121	
Application	Procedure	Last Modification Date	Last Modification Time
baseapp	flow01	2012/01/04	11.40.41
	flow01_dmhr_t	2012/01/04	11.40.41
	opt_test1	2014/03/25	11.57.34
	profile	2010/08/11	11.21.53
dimensions	ed_subtype_lookup	2016/05/23	15.45.56
	ed_subtype_lookup	2016/06/28	15.57.55
ibisamp	cargraph	2012/02/06	22.05.52
	carinst	2012/02/06	22.05.52
	carinst2	2012/02/06	22.05.52

ACI Monitoring

The ACI Monitoring report provides information about changes to the server configuration files, and user application files. Actual changes made are not reported, but who, when, and the type of access is shown. To access this report, you must have the Access Control setting that can be found in the Repository Properties pane enabled.

You can specify the following parameters in the Report Management report filter page for this report:

- User name.** Allows you to run the report based on user. The default value is All.
- Option.** Allows you to filter based on the type of access. Options are All, Read file, Write file, and Delete file. The default option is All.
- Category.** Allows you to specify the types of files that you want to audit. Options are All, System files, or User files. The default value is All.
- E-Mail.** Allows you to view the sender and receiver information in the form of their individual email address for any email messages that were sent in the Web Console.

The following image shows the ACI Monitoring report.

ACI - Summary Report
 Category: **ALL**
 Server Name: FUENT-PC:8121
 Date Range: 09/06/2018 - 09/13/2018

Date	User	Time	Action	Connection Address	File Type	Source Application/Filename	Target Application/Filename	Return Code	Status
09/13/2018	IBI\cf12026	13:13:05.839	READ	:::1	prf	EDASPROF.prf			
		13:13:16.147	WRITE	:::1	prf	edasprof.prf			
		13:13:17.463	READ	:::1	prf	EDASPROF.prf			
		13:15:50.180	WRITE	:::1	fex	ibisamp/prc01.fex			
		13:16:10.015	READ	172.30.240.58	cfg	_edaconfcatalog/irm/gktable.cfg			
		13:16:57.732	APP COPY	172.30.240.58	Application	A		224	Error
		13:18:59.269	READ	172.30.240.58	fex	ibisamp/prc01			

This report has two hyperlinks that allow you to drill down to additional information. The following table lists and describes these hyperlinks.

Click a hyperlink in the:	Description
Action column	<p>Opens a dialog box that shows the actual command that was issued for the operation.</p> <p>If you filter a report by <i>E-Mail</i>, the hyperlink will open a dialog box that shows the email address for both the sender and receiver of an email sent in the Web Console.</p>
Return Code column	<p>Opens a dialog box that provides further explanation on the error code listed.</p>

Resource Governor Reports

The following section provides a description of each report in the Resource Governor folder.

To access these reports in the Web Console, expand the *Reports* folder in the tree. Then, expand the *Resource Governor* folder, right-click the desired report, and select *Run*.

Monitored Data Sources

The Monitored Data Sources report provides an overview of the data sources that are being monitored, including number of requests, and the first and last access dates.

To access the Monitored Data Sources report in the Web Console, click Resource Management on the toolbar and expand the Reports folder. Right-click *Monitored Data Sources*, and click *Run* from the context menu. After specifying report filters, click *View Report*.

The following image shows the Monitored Data Sources report.

Monitored Data Sources									
Server Name: FUENT-PC:8121									
Date Range: 08/12/2016 - 10/27/2016									
Application Name	Data Source Name	Number of Requests	Duration	CPU Time	Wait Time	Adapter Time	I/O Used	Records Processed	Rows Returned
ibisamp	car	1	00:00:00.0620	00:00:00.0150	00:00:00.0470	00:00:00.0000	4	18	18
Total For Application: ibisamp		1	00:00:00.0620	00:00:00.0150	00:00:00.0470	00:00:00.0000	4	18	18
ibisamp/dimensions	wf_retail_product	3	00:00:00.1100	00:00:00.0460	00:00:00.0640	00:00:00.0690	161	477	161
Total For Application: ibisamp/dimensions		3	00:00:00.1100	00:00:00.0460	00:00:00.0640	00:00:00.0690	161	477	161

Monitored Procedures

The Monitored Procedures report provides an overview of the procedures that are being monitored, including the number of executions, and the first and last execution dates.

To access the Monitored Procedures report in the Web Console, click Resource Management on the toolbar and expand the Resource Governor folder. Right-click *Monitored Procedures*, and click *Run* from the context menu. After specifying report filters, click *View Report*.

The following image shows the Monitored Procedures report.

Monitored Procedures									
Server Name: FUENT-PC:8121									
Date Range: 08/12/2016 - 10/27/2016									
Application Name	Procedure Name	Number of Executions	Duration	CPU Time	Wait Time	Adapter Time	I/O Used	Records Processed	Rows Returned
baseapp	prc01	1	00:00:00.0060	00:00:00.0060	00:00:00.0000	00:00:00.0000	1	0	0
	ra_test1	1	00:00:00.0440	00:00:00.0000	00:00:00.0440	00:00:00.0000	0	0	0
	sales_by_brand	1	00:00:00.1450	00:00:00.0000	00:00:00.1450	00:00:00.1230	0	24	24
	testerror	1	00:00:00.0630	00:00:00.0000	00:00:00.0630	00:00:00.0000	1	0	0
Total For Application: baseapp		4	00:00:00.2580	00:00:00.0060	00:00:00.2520	00:00:00.1230	2	24	24
ibisamp	carinst	1	00:00:00.0680	00:00:00.0150	00:00:00.0530	00:00:00.0000	4	18	18

Resource Thresholds Report

The Threshold Adviser Report shows parameter details of all Resource Governor thresholds that have been built.

To access the Resource Thresholds report in the Web Console, click Resource Management on the toolbar and expand the Resource Governor folder. Right-click *Resource Thresholds*, and click *Run* from the context menu. After specifying report filters, click *View Report*.

The following image shows the Resource Thresholds report.

Resource Thresholds													
Server Name: FUENT-PC:8121													
Threshold Name	Threshold Type	Execution Time	CPU Time	IO Used	Rows Returned	Active	Start Date	End Date	Start Time	End Time	Monday	Tuesday	Wednesday
DEFAULT	Date and Time	1440	3600	100000	100000	Yes	0101	1231	0000	2359	Off	Off	Off
RG	Daily	1440	3600	100000	100000	Yes	0101	1231	0000	2359	Off	Off	On

Knowledgebase Properties Report

The Knowledgebase Properties Report shows details of all knowledge bases that have been built, the rule file used, and the status of the knowledge base.

To access the Knowledgebase Properties report in the Web Console, click Resource Management on the toolbar and expand the Resource Governor folder. Right-click *Knowledgebase Properties*, and click *Run* from the context menu. After specifying report filters, click *View Report*.

The following image shows the Knowledgebase Properties report.

Knowledgebase Properties							
Server Name: FUENT-PC:8121							
Date Range: 01/01/1995 - 10/27/2016							
Knowledgebase Name	Knowledgebase Type	Rules File Name	Rules File Status	Govern Status	Date Created	Data Source Name	Monitored Status
CAR_RULE1	Business	GKECR	Used	On	10/25/2016	ibisamp/car	Monitored
CAR_RULE2	Data	-		Off	10/26/2016	ibisamp/car	Monitored
RG	Global	BR1	Used	Off	10/26/2016	N/A	N/A
RGBR1	Global	BR1	Used	Off	10/25/2016	N/A	N/A

This report has two hyperlinks that allows you to drill down to other reports, as described in the following table.

Click a hyperlink in the:	Report Generated	Description
Knowledgebase Name column	Knowledgebase Lines	

Click a hyperlink in the:	Report Generated	Description
Rules File Name column	The rules associated to the data sources.	

Threshold Advisor Report

The Threshold Advisor Report recommends thresholds based on statistical information derived from monitored data.

To access the Threshold Advisor report in the Web Console, click Resource Management on the toolbar and expand the Resource Governor folder. Right-click *Threshold Advisor*, and click *Run* from the context menu. After specifying report filters, click *View Report*.

The following image shows the Threshold Advisor report.

Threshold Advisor
Recommended thresholds for: **ALL**
Server Name: **FUENT-PC:8121**

Application Name	Data Source Name	Data Type	Execution Time Threshold	CPU Time Threshold	Rows Threshold	IOs Threshold
ibisamp	car	EDA	0	0	18	4
ibisamp/dimensions	wf_retail_product	EDA	0	0	159	159
unknown	hold1	EDA	0	0	286	286
	wfrhold1	EDA	0	0	20	62
	wfrhold10	EDA	0	0	20	62

Governed Requests Report

This report shows which requests been governed or canceled and the rule file used to govern.

Governing Reports can be generated for particular data sources and user IDs, for all data sources and user IDs, or for combination of the two.

To access these reports in the Web Console, expand the *Resource Governor* folder in the tree, right-click *Governed Requests* and select *Run*. The Report filter for Governed Requests window opens, as shown in the following image. Here you can filter on data source name and/or user. Enter a valid data source name and/or user name to filter on, or enter ALL for all data sources and all users. The default is ALL.

Report filter for Governed Requests

Select server: Not applicable, this report only includes local data

Report format: HTML

Filter on Data Source name and/or User name selected.

Data Source name: All (Select a Data Source name or All)

User name: All (Select a User name or All) Report period: Date range

Start Date: 📅 2019/2/19

End Date: 📅 2019/2/22

View Report

The following image shows an example of the Governed Requests report.

Governed Requests												
Server Name: FUENT-PC:8121												
Date Range: 08/12/2016 - 10/28/2016												
Date	Time	User	Application Name	Data Source Name	Data Type	Knowledgebase Name	Status	Rule Number	Threshold Exceeded	Reason	Rule File Type	Procedure Name
10/28/2016	11:46:31	IBI\CF12026	ibisamp	car	EDA	GBR1	Canceled	0001		Select *	Global	Ad-Hoc

Getting Help

The following section explains the help features of Resource Management in the Web Console.

In this chapter:

- [Resource Management Administration Help](#)
 - [Reporting Application Context-Sensitive Help From Reports](#)
-

Resource Management Administration Help

Online help for Resource Management is available from the Web Console.

Web Console Help

To access online help from the Web Console, click the *Help* button located in the upper right corner of the console window and click *Contents and Search*, as shown in the following image.

The Web Console server help page opens.

The following image shows the Web Console server page, which displays several available help topics, including two administration topics: *Resource Analyzer and Resource Governor* and *Resource Analyzer Report Options*. The first topic is selected in this image.

The screenshot displays the Web Console help interface. On the left is a navigation pane with a search bar and a list of help topics. The 'Resource Analyzer and Resource Governor' section is expanded, and 'Resource Management Interface' is selected. The main content area shows the title 'Resource Analyzer and Resource Governor' with an 'INDEX' button. Below the title, a 'Topics:' section lists three items: 'Resource Management Interface', 'Administering Resource Management From the Web Console', and 'Migrating the Resource Management Repository'. The first item is selected. To the right of this list is a detailed text description of the Resource Analyzer and Resource Governor components. Below the text is a 'Top of page' link. A second section titled 'Resource Management Interface' contains an 'In this section:' list with 'Ribbon' and 'Ribbon Organization'. At the bottom, there is explanatory text about the ribbon functionality.

Search for

All Documents

Max. 5 matches per document

Web Console Help

- Server Introduction
- Server Security
- Managing Applications
- Managing Adapters
- Metadata
- Stored Procedures
- Server Workspace Manager
- Listeners and Special Services
- Troubleshooting
- Unicode Support
- Platform-Specific Methods for Specifying EDAEXTSEC
- Resource Analyzer and Resource Governor**
 - Resource Management Interface**
 - Administering Resource Management From the Web Console
 - Migrating the Resource Management Repository
 - Resource Analyzer Report Options
- Setting Up DataMigrator
- Adapters
- Reference
- Index

Resource Analyzer and Resource Governor INDEX

Topics:

- Resource Management Interface**
- Administering Resource Management From the Web Console
- Migrating the Resource Management Repository

Resource Analyzer and Resource Governor are optional components of the server. Resource Analyzer monitors data usage activity by collecting the attributes of requests and storing them in a set of Usage Monitoring tables. It automatically keeps track of each request, which data sources and columns it accessed, when it was run, how long it took, which resources it consumed, and more. Resource Analyzer allows you to report on, graph, and analyze end-user request traffic.

Resource Governor controls monitoring, system configuration parameters, and governing rules. It provides preemptive governing for requests issued to both relational and non-relational data sources and uses the same Usage Monitoring tables as Resource Analyzer. An administrator can also define site-specific thresholds controlling the amount of resources a request can use. Based on the usage and threshold information, Resource Governor builds rules about how to govern requests against specific data sources. The governing facility uses the rules when inspecting each request, stopping any request that is estimated to exceed the predetermined resource thresholds.

[Top of page](#)

Resource Management Interface

In this section:

- Ribbon
- Ribbon Organization


The Resource Management Web Console functionality is available through a ribbon at the top of the interface. The Resource Management ribbon replaces traditional toolbars with a single, rectangular region that spans the width of the application window.

There are two ribbons in Resource Management. The following image shows the ribbon before configuration has taken place.

Reporting Application Context-Sensitive Help From Reports

Resource Management reports offer context-sensitive online help from almost every report. To access the online help for a specific report, run a report and then click the *Help* link that appears at the top of the report.

When you click the *Help* link in a report, a new browser window opens to display a topic relevant to that report. The window also displays a table of contents that allows you to easily navigate to help topics for other reports. Simply click a topic on the left, and the information displays on the right. For example, when you click the *Help* link while viewing the Monitored Sessions report, the window displays the Monitored Sessions topic, as shown in the following image.



Administrative Usage Monitoring Tables Column Descriptions

This appendix provides descriptions of the Administrative and Usage Monitoring tables. Each description shows the table column definitions.

In this appendix:

- [Usage Monitoring Table \(RMLDB.MAS\)](#)
 - [Usage Monitoring RDBMS Table Sizing](#)
 - [Renamed or Moved Columns](#)
-

Usage Monitoring Table (RMLDB.MAS)

This section lists the data definitions that comprise the Resource Analyzer Usage Monitoring Table, and provides an explanation of the segments and field values.

Note:

- The CPU, Elapsed, Wait, DBMS, and zIIP time values in all segments are in 10,000th of a second increments and are indicated by a double asterisk (**) next to the field name.
- A hash tag (#) next to a field denotes that it is new in this release. A dollar sign (\$) next to a field denotes that it has been renamed. A DEFINE has been added to the Master File Description that redefines the new field name to the old field name for compatibility purposes.
- On MVS, for both HFS and PDS deployments, the MVSAPP entry in the servers edasprof profile should not be removed. If it is, some data field values may not be collected and will be blank.

SMSERVERS Segment

This segment keeps track of server related information. It is the root segment record and it is updated every time the database is updated with new log information and when system maintenance is performed to remove old data.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMSRVNAMES	Alphanumeric, length=128V	SERVERNAME value.
SITECODE	Alphanumeric, length=7	Information Builders site code value.
MINSESSDATE	Alphanumeric, length=8	The first session date in the database for this server. The format is YYYYMMDD.
MAXSESSDATE	Alphanumeric, length=8	The last session date in the database for this server. The format is YYYYMMDD.
MONSTATUS	Alphanumeric, length=1	Monitoring status: I=Inactive, A=Active.
CNFGDATE	Alphanumeric, length=8	Configuration date.
LAST_MAINT	Alphanumeric, length=8	Date of last repository maintenance.
LAST_UPDATE	Alphanumeric, length=8	Date of last repository update.
SESSCOUNT	Double	Count of sessions in repository.
DBRELEASE	Alphanumeric, length=10	Repository release number.
SRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMSESS Segment

This segment contains a single record for each connected session for the server. It is the parent segment for the procedure executed in a session and for requests executed in a session. This record contains the accumulated values for all activity in a connected session.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SESSKEY	Alphanumeric, length=38	Session key. This key is a unique value for each session for a server.
SMSERVER	Alphanumeric, length=128V	The server the session ran on. This is a required field to support relational databases used by the cluster master.
SESSDATE	Alphanumeric, length=8	The session date. The format is YYYYMMDD.
SESSSTART	Alphanumeric, length=9	The session start time. The format is HHMMSSNNN.
SESSSTOP	Alphanumeric, length=9	The session end time. The format is HHMMSSNNN.
SESSDUR**	Double	The session duration time.
SESSCPU**	Double	The session CPU time.
SESSWAIT**	Double	The session wait time.
SESSIOS	Double	The number of I/O operations used in the session.
SESSRECORDS	Double	The number of RECORDS processed in the session.
SESSTRANS	Double	The number of TRANSACTIONS processed in the session.
SESSLINES	Double	The number of LINES returned or held in the session.
SMUSERID	Alphanumeric, length=128V	The effective user ID for the session. This is the connection user ID unless overridden.
USERIDTYPE	Alphanumeric, length=1	GKE %(M)RE USERID GKE %(U)SERID override value. Values are M, U, or blank.

Field	Value	Description
SMLIVE	Alphanumeric, length=1	Indicates the type of session. Values are 0, 1, 2, 3, 4, or 5. See the SESSTYPE definition for descriptive values.
SESSFEXCNT	Numeric, length=4	Number of FOCEXECs monitored during the session.
SESSCMDCNT	Numeric, length=4	Number of commands monitored throughout the session.
SESSAUDCNT#	Numeric, length=4	The number of SMAUDIT segment records for this session.
SESSPREF	Alphanumeric, length=64	Monitor preference setting when session started.
SMCONNTYPE	Alphanumeric, length=1	See NETWORK DEFINE.
SMCONNID	Alphanumeric, length=128V	Connection user ID.
SMGROUP	Alphanumeric, length=128V	Security group for SMCONNID.
SMCONNADDR	Alphanumeric, length=48	Network connection address.
CONNRETCODE	Alphanumeric, length=10	Connection return code.
SMCLIENTTYPE	Alphanumeric, length=32	CLIENT value if sent from the client.
SMFOCREL	Alphanumeric, length=30	Usage Monitor release and build number.
SMSEZIIPTIME**	Double	The zIIP processor time used. This value is only available on MF processors running z/OS.

Field	Value	Description
SMSEZIIPONCP**	Double	The zIIP on CP value. This value is only available on MF processors running z/OS.
SMSEZIIPON	Alphanumeric, length=1	Indicates if zIIP was enabled. Values are Y, N, or blank. This value is only available on MF processors running z/OS.
SESSCPUID	Alphanumeric, length=48	CPUID value.
SESSLPARNM	Alphanumeric, length=8	The LPAR name the session ran on. This value is only available on MF processors running z/OS.
SESSMODELID	Alphanumeric, length=8	The CPU model ID the session ran on. This value is only available on MF processors running z/OS.
SMCLNTCODEPG	Double	The code page of the NLS client.
SMJOBID	Alphanumeric, length=63	Job ID of the scheduled or deferred job.
SMFOCUSIO	Double	The number of FOCUS I/O operations in the session.
SMDBMSIO	Double	The number of RDBMS adapter I/O operations used in the session.
SMDBMSTIME**	Double	Elapsed time spent in RDBMS.
SMMEMUSAGE	Double	Megabytes of memory used in the session. This value includes Fixed Memory on MF processors running z/OS.
SMDISKUSAGE	Double	Kilobytes of disk space used in the session.
SMSVCNAME	Alphanumeric, length=64	Agent Service Name.
SESSSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMMEMORY	Double	Kilobytes(KB) of swappable memory used.

Field	Value	Description
SMMEMFIXED	Double	Kilobytes(KB) of fixed memory used. This value is only available on MF processors running z/OS.
SMSYSPCTCPUUSAGE#	Numeric, length=4	The percent of system CPU usage at session start.
SMSYSMEMORY#	Double	The amount of memory available at session start.
SMSYSFREEMEMORY#	Double	The amount of free memory available at session start.
SMSYSVIRTMEMORY#	Double	The amount of virtual memory available at session start.
SMSYSFREEVIRTMEMORY#	Double	The amount of free virtual memory available at session start.

SMRPCS Segment

This segment contains a single record for each executed or included procedure. It is the parent segment for SMRPCCMD and SMRPCWF segments. The resource values in the segment are a subset of the resource values of its parent segment, and it contains the resource values used by the FOCUS or SQL requests it may execute.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMRPCKEY	Alphanumeric, length=38	Session key value.
SMRPCNUM	Numeric, length=4	The execution order number. Values are 1 to 2,147,483,647.
SMRPCDATE	Alphanumeric, length=8	The date the procedure was executed. The format is YYYYMMDD.
SMRPC TIME	Alphanumeric, length=9	The time the procedure was started. The format is HHMMSSNNN.

Field	Value	Description
SMRPCSTOP	Alphanumeric, length=9	The time the procedure ended. The format is HHMMSSNNN.
SMFEXNAME	Alphanumeric, length=577V	The procedure name. This is the WebFOCUS physical name. The path to the file is located in SMRPCWF.SMWFBASEDIR
SMRPCETIME**	Double	The procedure execution time.
SMRPCCTIME**	Double	The procedure CPU time.
SMRPCWTIME**	Double	The procedure wait time.
SMRPCIOS	Double	The number of I/O operations by the procedure.
SMRPCRECORDS	Double	The number of RECORDS processed by the procedure.
SMRPCTRANS	Double	The number of TRANSACTIONS processed by the procedure.
SMRPCCLINES	Double	The number of LINES returned or held by the procedure.
SMRPCPNUM	Alphanumeric, length=4	Parent procedure number or 0.
SMRPCATYPE	Alphanumeric, length=1	Type of procedure. I (-INCLUDE), E (EX), M (MFD_profile), or S (Stored Procedure Call).
SMRPCEND	Alphanumeric, length=1	Indicates how the procedure ended. Values are E, Q, or blank. E = Exit or Run, Q = Quit.
SMRPCWF	Alphanumeric, length=1	Request originated from WebFOCUS.
SMRPCWC	Alphanumeric, length=1	Request originated from Web Console.

Field	Value	Description
SMRPCLOC	Alphanumeric, length=254V	Physical location of the procedure.
SMRPCZIIPTIME**	Double	The zIIP processor time used. This value is only available on MF processors running z/OS.
SMRPCZIIPONCP**	Double	The zIIP on CP values. This value is only available on MF processors running z/OS.
SMRPCZIIPON	Alphanumeric, length=1	Indicates if zIIP was enabled. Values are Y, N, or blank. This value is only available on MF processors running z/OS.
SMRPCFOCUSIO	Double	The number of FOCUS I/O operations by the procedure.
SMRPCDBMSIO	Double	The number of RDBMS adapter I/O operations by the procedure.
SMRPCDBMSTIME**	Double	Elapsed time spent in RDBMS.
SMRPCSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMRPCPCTCPUUSAGE#	Numeric, length=4	The percent of system CPU usage at procedure start.
SMRPCSYSTEMMEMORY#	Double	The amount of memory available at procedure start.
SMRPCSYSFREEMEMORY#	Double	The amount of free memory available at procedure start.
SMRPCSYSVIRTMEMORY#	Double	The amount of virtual memory available at procedure start.
SMRPCSYSFREEVIRTMEMORY#	Double	The amount of free virtual memory available at procedure start.
SMRPCNUMPARMS#	Numeric, length=4	The number of SMRPCCCMD segment records for this procedure.

SMRPCCMD Segment

This segment contains one or more records that comprise the execution statement with all parameter names and their values that were supplied.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMPPARMKEY\$	Alphanumeric, length=38	Session key value.
SMPPARMNUM\$	Numeric, length=4	Procedure number.
SMPPARMSEQ	Numeric, length=4	Procedure line segment number.
SMRPPARMS	Alphanumeric, length=1024V	Procedure parameter and value (1 pair per record).
SMRPCCSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMRPCWF Segment

This segment contains a single record that contains information sent to the Reporting Server from WebFOCUS or ReportCaster.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMRWFKEY\$	Alphanumeric, length=38	Session key value.
SMRWFNUM\$	Numeric, length=4	Procedure number.
SMWFFEXNAME	Alphanumeric, length=577V	WebFOCUS procedure name. Also DEFINE as SMWFRPCNAME.
SMWFBASEDIR	Alphanumeric, length=254V	WebFOCUS base directory path value. Also DEFINE as SMBASEDIR.

Field	Value	Description
SMWFDOMAIN	Alphanumeric, length=254V	WebFOCUS MR domain name value. Also DEFINE as SMDOMAIN.
SMWFAPPPATH	Alphanumeric, length=254V	WebFOCUS MR application directory value. Also DEFINE as SMAPPPATH.
SMRCSCHEDID	Alphanumeric, length=66	ReportCaster scheduler ID.
SMRCPROCID	Alphanumeric, length=66	ReportCaster process ID.
SMWFMRGROUP	Alphanumeric, length=A128V	Reserved for a future release.
SMWFSRVNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMQUERY Segment

This segment contains a record for each monitored request executed. It is the parent segment for SMCMDLN, SMRMTS, SMFROMS, SMFNCTNS and SMRELTNS segments. The resource values in the segment are a subset of the resource values of its parent segment.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMKEY	Alphanumeric, length=38	Session key value.
SMCMDNUM	Numeric, length=4	Request number.
SMCMDNAME	Alphanumeric, length=12	The name of the FOCUS or SQL request type.
SMDATE	Alphanumeric, length=8	The date the command was executed. Format is YYYYMMDD.

Field	Value	Description
SMTIME	Alphanumeric, length=9	The time the command started. The format is HHMMSSNNN.
SMSTOP	Alphanumeric, length=9	The time the command ended. The format is HHMMSSNNN.
SMELOPTIME**	Double	The request execution time.
SMCPUTIME**	Double	The request CPU time.
SMWAITTIME**	Double	The request wait time.
SMIOS	Double	The number of I/O operations used by the request.
SMRECORDS	Double	The number of RECORDS processed by the request.
SMTRANS	Double	The number of TRANSACTIONS processed by the request.
SMROWS	Double	The number of LINES returned or held by the request.
SMROWLEN	Numeric, length=4	Longest record length of rows held/returned.
SMOUTLIMIT	Numeric, length=4	SQL and FOCUS request max output limit.
SMRECLIMIT	Numeric, length=4	FOCUS requests of max RECORDS read. SQL requests of max OUTPUT records.
SMREADLIMIT	Numeric, length=4	Max READS for MODIFY requests.
SMRETCODE	Numeric, length=4	Command return code.
SMUNIONS	Numeric, length=4	Number of SQL UNION statements.
SMUNIONALLS	Numeric, length=4	Number of SQL UNION ALL statements.
SMNUMFROMS	Numeric, length=4	Number of data sources used.

Field	Value	Description
SMNUMCOLUMNS	Numeric, length=4	Number of columns used.
SMNUMRELATIONS	Numeric, length=4	Number of relational clauses.
SMNUMGROUPBYS	Numeric, length=4	Number of GROUP BYs.
SMNUMORDERBYS	Numeric, length=4	Number of SQL ORDER BY or FOCUS BY.
SMNUMFUNCTIONS	Numeric, length=4	Number of functions used.
SMNUMREMOTES	Numeric, length=4	Number of remote connections used.
SMSORTIO	Double	Number of sort I/O operations used.
SMEXTSORT	Alphanumeric, length=1	Indicates whether an external sort was used. Values are Y or N.
SMREQTYPE	Alphanumeric, length=1	Type of request. See REQUESTTYPE DEFINE in the Master File.
SMMDIFYTYPE	Alphanumeric, length=1	Type of MODIFY or MAINTAIN. See MODIFYTYPE DEFINE in the Master File.
SMCOLLECT	Alphanumeric, length=1	Reason for collection. See SMSTATUS DEFINE in the Master File.
SMALLROWS	Alphanumeric, length=1	Indicates whether all rows were returned: 0 = Incomplete. 1 = Complete. 2 = Data adapter limited.
SMALLCOLS	Alphanumeric, length=1	Indicates whether a SELECT * or PRINT * used. Values are Y or N.
SMSETALL	Alphanumeric, length=1	Indicates whether SET ALL was used: 0 = OFF 1 = ON 2 = PASS

Field	Value	Description
SMCORRQRY	Alphanumeric, length=1	Indicates whether the SQL request was correlated. Values are Y, N, or blank.
SMHLDTYPE	Alphanumeric, length=1	Type of HOLD file. See HOLDTYPE DEFINE in the Master File.
SMHLDFORMAT	Alphanumeric, length=2	Format of hold file. See HOLDFORMAT DEFINE in the Master File.
SMHLDNAME	Alphanumeric, length=577V	HOLD file name.
SMRPCNAME	Alphanumeric, length=577V	Procedure name the request originated from. Blanks = Ad hoc.
SMFEXLNO	Alphanumeric, length=10	Procedure line number if SMRPCNAME is not blank.
SMFEXLOC	Alphanumeric, length=254V	Physical location of the procedure if SMRPCNAME is not blank.
SMZIIPTIME**	Double	The zIIP processor time used. This value is only available on MF processors running z/OS.
SMZIIPONCP**	Double	The zIIP on CP value. This value is only available on MF processors running z/OS.
SMZIIPON	Alphanumeric, length=1	Indicates if zIIP was enabled. Values are Y, N, or blank. This value is only available on MF processors running z/OS.
SMCMDFOCUSIO	Double	The number of FOCUS I/O operations by the request.
SMCMDDBMSIO	Double	The number of RDBMS adapter I/O operations by the request.
SMRPCDBMSTIME**	Double	Elapsed time spent in RDBMS.
SMQPRIMRPCNUM	Numeric, length=4	SMRPCS.SMRPCNUM value of the primary procedure or 0.

Field	Value	Description
SMQRPCNUM	Numeric, length=4	SMRPCS.SMRPCNUM value or 0.
SMREMARKS	Alphanumeric, length=254V	Notes from admins, updated by GKE %REMARK command.
SMQSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMNUMCMDLNS#	Numeric, length=4	The number of SMCMDLN segment records for this request.

SMCMDLN Segment

This segment contains one or more records that comprise the original request.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMCLNKEY	Alphanumeric, length=38	Session key value.
SMCLNCNUM\$	Numeric, length=4	Command number.
SMCLNSEQ\$	Numeric, length=4	Command line segment number.
SMCMDLINE	Alphanumeric, length=72V	Command line.
SMCLSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMGOVEND Segment

This segment is updated if the Governor cancels the request and issues a Cancel message or issues an Advise message.

The following table lists the available fields, their values, and descriptions.

Note: Field values in the SMGOVEND segment with an exclamation point (!) after the name indicate the length value for the field has changed from 8 to 48. Due to limitations of the operating system, the PDS deployed servers will only allow a maximum of eight characters to be stored in the field.

Field	Value	Description
SMGOVKEY	Alphanumeric, length=38	The session key.
SMGOVCNUM\$	Numeric, length=4	Request number.
SMKBNAME	Alphanumeric, length=48	Compiled rule file used.
SMRULENUM	Alphanumeric, length=4	Rule number that determined the result or decision.
SMTHRESHTYPE	Alphanumeric, length=8	Type of threshold exceeded.
SMTHRESHHOLD	Alphanumeric, length=10	Value of threshold exceeded.
SMGDBMS	Alphanumeric, length=8	EDA or SQL engine suffix.
SMGKNBTYPE	Alphanumeric, length=6	Type of compiled rule file. Values are Both, Data, Custom, or Global.
SMGDSNAME	Alphanumeric, length=577V	Data source for compiled rule file.
SMGOVNTYPE	Alphanumeric, length=1	C = Canceled, A = Advised
SMGSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMRMTS Segment

This segment contains the remote statistics for a request that retrieves data from a subserver. The resource values of this segment are in addition to the resource values of its parent segment.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMRMTKEY	Alphanumeric, length=38	The session key.
SMRMTCNUM\$	Numeric, length=4	The request number.
SMRMTSEQ\$	Numeric, length=4	Remote number.
SMRMTETIME	Double	Remote session duration time.
SMRMTCTIME	Double	Remote session CPU time.
SMRMTIOS	Double	The number of I/O operations used in the remote session.
SMRMTROWS	Double	The number of rows returned from the remote session.
SMRMTSQLCODE	Alphanumeric, length=11	The SQL return code from the remote request.
SMRMTPID	Alphanumeric, length=11	Process ID of the session on the remote server.
SMRMTDATE	Alphanumeric, length=8	The local date on the remote server.
SMRMTTIME	Alphanumeric, length=9	Local time on the remote server.
SMRMTNAME	Alphanumeric, length=128V	Remote server name.
SMRMTSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMFNCTNS Segment

This segment contains a record for each function used in a monitored request.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMFUNKEY	Alphanumeric, length=38	The session key.
SMFUNCNUM\$	Numeric, length=4	The request number.
SMFUNSEQ\$	Numeric, length=4	The function number.
SMFUNCTION	Alphanumeric, length=12	The function name.
SMFNAME	Alphanumeric, length=577V	The data source that the field used.
SMFCOLUMN	Alphanumeric, length=512V	The field that the function used.
SMASTKUSED	Alphanumeric, length=1	Indicates whether or not an asterisk was used by the function. Values are Y or N.
SMFLITERAL	Alphanumeric, length=32	The literal value, if used.
SMFSRVNUM	Numeric, length=4	Used in many-to-one configurations.

SMRELTNS Segment

This segment contains a record for each WHERE or IF clause used in a monitored request.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMRELKEY	Alphanumeric, length=38	The session key.
SMRELCNUM\$	Numeric, length=4	The request number.
SMRELSEQ\$	Numeric, length=4	The relation number.

Field	Value	Description
SMLNAME	Alphanumeric, length=577V	Data source name used in the LHS.
SMLCOLUMN	Alphanumeric, length=512V	Field name used in the LHS.
SMRNAME	Alphanumeric, length=577V	Data source name used in the RHS, or blanks.
SMRCOLUMN	Alphanumeric, length=512V	Field name used in the RHS, or blanks.
SMPERATOR	Alphanumeric, length=8	Operator used in the clause. For example, <, =, !=, etc.
SMANDOR	Alphanumeric, length=3	The AND or OR relational connector used.
SMNOT	Alphanumeric, length=1	The NOT keyword used. Values are Y or N.
SMALL	Alphanumeric, length=1	The ALL keyword used. Values are Y or N.
SMANY	Alphanumeric, length=1	The ANY keyword used. Values are Y or N.
SMLITERAL	Alphanumeric, length=32	The literal value, if used.
SMRSRVNUM	Numeric, length=4	Server number. Used in many-to-one configurations.

SMFROMS Segment

This segment contains a record for each data source used in a monitored request. It is the parent segment for the SMCOLMNS segments.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMFRMKEY	Alphanumeric, length=38	The session key.
SMFRMCNUM\$	Numeric, length=4	The request number.
SMDSNUM	Numeric, length=4	The data source number.
SMNAME	Alphanumeric, length=577V	The data source name.
SMSUFFIX	Alphanumeric, length=8	The data source suffix, or blanks.
SMDBMS	Alphanumeric, length=8	The EDA or SQL engine suffix.
SMDATATYPE	Alphanumeric, length=1	The data source type: P = Permanent T = Temporary R = Relational
SMMORE	Alphanumeric, length=1	Used as MORE FILE. Values are Y, N, or blank.
SMSUNAME	Alphanumeric, length=8	The SU identifier.
SMALIAS	Alphanumeric, length=128V	The USE AS alias name.
SMDBLOC	Alphanumeric, length=577V	The data source location. Formerly named SMDBNAME.
SMMASLOC	Alphanumeric, length=577V	The Master File location.
SMTBLNAME	Alphanumeric, length=577V	Table name of the relational table.
SMFRSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMNUMCOLS#	Numeric, length=4	The number SMCOLMNS segment records for this Data Source.

SMCOLMNS Segment

This segment contains a record for each field retrieved from a data source in a monitored request.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMCOLKEY	Alphanumeric, length=38	The session key.
SMCOLCNUM\$	Numeric, length=4	The request number.
SMCOLDSNUM\$	Numeric, length=4	The data source number.
SMCOLSEQ\$	Numeric, length=4	The column number.
SMCOLUMN	Alphanumeric, length=512V	The column/field name.
SMDEFINE	Alphanumeric, length=1	Column type: 0 = UNKNOWN 1 = REAL 2 = PERMANENT DEFINE 3 = TEMPORARY DEFINE 4 = COMPUTE
SMIDXKEY	Alphanumeric, length=1	The indexed field. Values are Y or N.
SMBYTYPE	Alphanumeric, length=1	BY type: 1 = ORDER BY or BY 2 = GROUP BY 3 = ORDER BY & GROUP BY, or blank. Not used with BY.

Field	Value	Description
SMBYORDER	Alphanumeric, length=1	Sort order: ascending (A), descending (D), or blank.
SMCOLSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMDVPROP	Alphanumeric, length=1	Dimension Property: Dimension (D), Measure (M), or blank.
SMASNAME#	Alphanumeric, length=512V	The AS value used in a request. This will only be populated if SET ASNAME ON is used before the request.

SMAUDIT Segment

This segment will contain the translated request if the original submitted request was transformed from SQL to FOCUS or from FOCUS to SQL. It will also contain any messages issued.

The following table lists the available fields, their values, and descriptions.

Field	Value	Description
SMAUDITKEY	Alphanumeric, length=38	Session key.
SMAUDITNUM	Numeric, length=4	Audit Sequence number.
SMAUDITFEXNUM	Numeric, length=4	Procedure number that caused this record to be created.
SMAUDITCMDNUM	Numeric, length=4	Request number that caused this record to be created.
SMAUDITTYPE	Alphanumeric, length=12	Audit record type. Values are: <input type="checkbox"/> Focus2SQL <input type="checkbox"/> Transform
SMAUDITREC	Alphanumeric, length=512V	Audit record.

Field	Value	Description
SMAUDITSRVRNUM	Numeric, length=4	Server number. Used in many-to-one configurations.
SMAUDITEMGSRV	Alphanumeric, length=1	Y (Message was displayed), N (Message was not displayed), or blank (Not a message record).

Usage Monitoring RDBMS Table Sizing

The recommended size of the relational database tables will vary from one implementation to another. The following are the major factors to consider:

- Type of configuration: Standard or Many-to-One.
- How many requests are anticipated from all configured servers?
- The level of monitoring that will be used on each server:
 - There are two levels of detail in Stored Procedure monitoring.
 - There are seven levels of detail in Command monitoring.

For more information on what levels of detail are available, see [Configuring and Administering Resource Management](#) on page 21. The sizing example below assumes that *all* levels are being monitored.

The storage values shown in the table below are estimates. Resource Management uses VARCHAR for character fields that are greater than or equal to 120 bytes in length, except for the SMREQUESTS and SMRPCREQUESTS tables, which use 72 byte VARCHAR fields. All VARCHAR fields are the ACTUAL size plus 2 bytes. This size difference is used in the Table Size chart but not in the tables used to describe each field.

Note: The table with a \$ is used by Resource Governor. It must be allocated, but the size can be reduced if you are not using Resource Governor.

Table Name	Number of Rows	Row Size	Table Size in KB*	Notes
SMSERVERS	1	220	1	1- <i>n</i> rows. Where <i>n</i> is the number of configured servers in the repository.
SMSESSIONS	30000	1151	59011	One row per session.

Table Name	Number of Rows	Row Size	Table Size in KB*	Notes
SMQUERY	300000	2277	1167407	One row per data request.
SMREQUESTS	3000000	124	635724	1- <i>n</i> rows per data request, where <i>n</i> is the number of 72 character segments in the data request.
SMGOVERN\$	15000	710	18175	One row per governed request.
SMREMOTES	20000	251	8579	1- <i>n</i> rows per suffix=EDA data request, where <i>n</i> is the number of subservers used in the request.
SMFROMS	750000	2526	3237671	One row for each data source used.
SMCOLUMNS	1500000	1087	2786499	One row per column returned or held.
SMFUNCTIONS	50000	1188	101514	One row per function used.
SMRELATIONS	150000	2282	584985	One row per relational test.
SMRPCS	35000	1375	82245	One row per stored procedure.
SMRPCREQUESTS	40000	1076	73555	One row per parameter pass to the procedure. For example, EX MYCAR PARM1=A, PARM2=B will create two rows. If the procedure is executed without any parameters, then no rows are created.
SMRPCWF	35000	2624	156953	One row per stored procedure from WebFOCUS.
SMAUDIT	3000000	581	2973633	1- <i>n</i> rows per translated request or message issued in a session.

Example of a WebFOCUS procedure:

```
EX MYWFFEX PAY=50000
```

Where MYWFFEX contains:

Renamed or Moved Columns

```
SELECT E.NAME, E.ADDRESS, E.JOBTITLE, J.JOBDESCRIPTION, J.PAY
FROM EMPLOYEE E, JOBS J
WHERE E.JOBTITLE = J.JOBTITLE AND J.PAY < &PAY
ORDER BY J.JOBTITLE, J.PAY;
```

Stores:

```
1 SMSESSION (1 *      1151)      1151
1 SMQUERY (1 * 2277)          2277
5 SMREQUESTS (4 * 124)         496
2 SMFROMS (2 * 2576)          5152
5 SMCOLUMNS (5 * 1087)       5435
2 SMRELATIONS (2 * 2282)      4564
1 SMRPCS (1 * 1375)           1375
1 SMRPCREQUESTS (1 * 1076)    1076
1 SMRPCWF (1 * 2624)          2624
-----
24129 bytes of storage
```

To determine the average storage needed, take the average of the least complex monitored request and the most complex, and multiply that value by the number of requests that could be monitored in a month or year.

The size of the table is calculated using a CP size, or character size, of 1. If the server is using a UNICODE code page (for example, CP65001), the size of the rows should be recalculated using a CP size of 3 to obtain the correct table size. The formula for calculating the table size is defined as:

$$(\text{Rows} * (\text{Size} * \text{CP size}) * 1.75) / 1024 = \text{KB}.$$

When creating a Db2 TABLESPACE, this value should be used for PRIQTY.

If the RDBMS being used for the repository uses a different byte size for nvarchar fields, use the value used by the RDBMS. A general rule is that nvarchar fields will use three bytes to represent a character. All table sizes have been rounded up as necessary.

Renamed or Moved Columns

This table identifies the columns that have been renamed or moved to new segments from their previous location in earlier releases of Resource Analyzer.

Previous Master	Previous Name	New Segment	New Name
SMRPCS	SMRPCNAME	SMRPCS	SMFEXNAME
SMRPCS	SMWFRPCNAME	SMRPCWF	SMWFRPCNAME

Previous Master	Previous Name	New Segment	New Name
SMRPCS	SMBASEDIR	SMRPCWF	SMWFBASEDIR
SMRPCS	SMFEXLOC	SMRPCS	SMRPCNAME
SMRPCS	SMDATE	SMRPCS	SMRPCDATE
SMRPCS	SMTIME	SMRPCS	SMRPCETIME
SMRPCS	SMELAPTIME	SMRPCS	SMRPCETIME
SMRPCS	SMCPUTIME	SMRPCS	SMRPCCTIME
SMRPCS	SMDOMAIN	SMRPCWF	SMWFDOMAIN
SMRPCS	SMAPPPDIR	SMRPCWF	SMWFAPPPDIR
SMRPCS	SMRCSCHEDID	SMRPCWF	SMRCSCHEDID
SMRPCS	SMRCPROCID	SMRPCWF	SMRCPROCID
SMRPCS	SMINCLUDED	SMRPCS	SMRPCINCL
SMQUERY	SMSERVER	SMSESS	SMSERVER
SMQUERY	SMUSERID	SMSESS	SMUSERID
SMQUERY	SMLIVE	SMSESS	SMLIVE
SMQUERY	SMCONNTYPE	SMSESS	SMCONNTYPE
SMQUERY	SMCONNID	SMSESS	SMCONNID
SMQUERY	SMGROUP	SMSESS	SMGROUP
SMQUERY	SMCONNADDR	SMSESS	SMCONNADDR
SMQUERY	SMCLIENTTYPE	SMSESS	SMCLIENTTYPE

Previous Master	Previous Name	New Segment	New Name
SMQUERY	SMFOCUSREL	SMSESS	SMFOCREL
SMQUERY	SMREFERERURL	SMSESS	SMREFERERURL
SMQUERY	SMRPCLNO	SMQUERY	SMFEXLNO
SMFNCTNS	SMNAME	SMFNCTNS	SMFNAME
SMFNCTNS	SMCOLUMN	SMFNCTNS	SMFCOLUMN
SMFNCTNS	SMLITERAL	SMFNCTNS	SMFLITERAL
SMFNCTNS	SMALLROWS	SMFNCTNS	SMASKUSED
SMBYS	SMBYTYPE	SMCOLMNS	SMBYTYPE

The SMSQLLINE field in the SMREQSTS master file from previous releases has been split, depending on its contents.

Monitored Commands (other than EX FOCEXEC)

Previous Master	Previous Name	New Segment	New Name
SMREQSTS	SMSQLLINE	SMCMDLN	SMCMDLINE

EX FOCEXEC

Previous Master	Previous Name	New Segment	New Name
SMREQSTS	SMSQLLINE	SMRPCCMD	SMRPCLINE

Time Fields in Resource Management

This appendix explains the time fields available in Resource Management.

In this appendix:

- [Time Field Definitions](#)
-

Time Field Definitions

All of the fields listed in this appendix can be found in the RMLDATA/RMLDB synonyms.

The following field contents are based on "wall clock" time and are stored in A9 fields in date/time (HHMMSSNNN) format. DEFINE fields are provided for individual components.

Session table - SMSESS

SESSSTART - SESSSTOP

These two fields cover the time a connection request causes the Resource Management (RM) code to be called and the time the connection is terminated.

RPC table - SMRPCS

SMRPCSTIME - SMRPCSTOP

These two fields cover the start and stop time of all the Remote Procedure Call (RPC) requests in a session. There is at least one RPC request per session but there can be more.

Command table - SMQUERY

SMTIME - SMSTOP

These two fields cover the start and stop time of the Command (for example, TABLE, TABLEF, and MODIFY). There can be one or more Commands processed per RPC per session.

The following fields are based on machine timers and the data is stored in 10,000th of a second. DEFINE fields are provided for times converted to seconds, Minutes, and Hours.

Session table - SMSESS

SESSDUR

This field contains the total duration of the session.

SESSCPU

This field contains the total CPU used during the session.

SMDBMSTIME

This field is the total of all SMCMDDBMSTIME values for all commands executed during the session.

RPC table - SMRPCS

SMRPCETIME

This field is the total elapsed time from the start of the RPC to the end of the RPC request.

SMRPCCTIME

This field is the total CPU time used from the start of the RPC to the end of the RPC request.

SMRPCDBMSTIME

This field is the total of all SMCMDDBMSTIME values for all commands executed by the RPC.

Command table - SMQUERY

SMELAPETIME

This field is the total elapsed time from the start of a command to the end of the command. This value contains the elapsed time of the command and also includes the elapsed time spent processing a command that uses DBMS time.

SMCPU TIME

This field is the total CPU time from the start of a command to the end of the command.

SMCMDDBMSTIME

This field is the total elapsed time from the start of the servers SQL Adapter processing to the end of the SQL adapters data retrieval.

This time field is only created at the request level and is the length of time spent by the various server processes that prepare the request, sends the request to the RDBMS server, and processes the answer set received. The DBMS time values for a SMRPCDBMSTIME and SMDBMSTIME field are accumulation values derived from all of the commands executed in the session.

When a SQL request that accesses data starts, the DBMS start timer is set to the current time in 10,000th of a second. When the SQL request ends, the DBMS end timer is set to the current time in 10,000th of a second. The DBMS time is a subset of time spent while the SQL request was being processed between the time Resource Analyzer started monitoring the request and the time Resource Analyzer finished. The CPU time consumed by the RDBMS while processing the request is not available to RM and is not included.

There are two different ways the DMBS timer is used:

- Direct Passthru Requests
- Synonym Requests


Direct Passthru Requests:

The SMCMDDBMSTIME timer starts when the DPT processor is called and ends when the DPT processor finishes processing. The length of time spent in this process is stored in the repository and is used in reporting as HH:MM:SS.nnnn. This type of request will have a small difference between the SMELAPTIME and the SMCMDDBMSTIME because of the minimal amount of time used by the server to prepare the request for processing. As stated before, SMCMDDBMSTIME is only an elapsed time value and includes both server SQL adapter time and RDBMS server time.

Synonym Requests:

As with Direct Passthru requests, the SMCMDDBMSTIME timer starts when the request type processor is called and ends when the processor finishes processing. The length of time spent in this process is stored in the repository and is used in reporting as HH:MM:SS.nnnn. This type of request will have a larger difference between the SMELAPTIME and the SMCMDDBMSTIME because of the additional processing time used to transform and optimize the request to SQL syntax that meets the requirements of the RDBMS server. The difference also includes the processing of the synonym before transformation and optimization and processing the returned answer set.

The following image shows an example of a session that executes one RPC that executes two commands.



The session duration and CPU timers are started when the connection to the server is established and ends when the session disconnects. At session start time, no DBMS timer is started. The RPC duration and CPU timers are started when the server calls the Resource Management RPC processor. Just like the session processing no DBMS timer is started. When an RPC executes a monitored command, the Resource Management command processor is called for initialization. After initialization finishes and Resource Management returns control to the server, one of two paths will be taken, either Direct Passthru or Synonym processing.

For an SQL command that uses a Direct Passthru command, the processor starts the server DBMS timer. When the answer set is retrieved, control is returned to the server and the DBMS timer is stopped. The Resource Management command processor is then called to end its processing.

For a synonym command that accesses data on a DBMS server the request is transformed and optimized, if needed, and then calls the SQL processor, which starts DBMS timer. When the answer set is retrieved, control is returned to the server and the DBMS timer is stopped. When the answer set has been formatted, Resource Management command process is called to end its processing.

During the call to end the Resource Management command process at the end of a request, the DBMS timer value of the server is copied to the SMCMDDBMSTIME field. It is also added to the SMRPCDBMSTIME and SMDBMSTIME fields. If the current RPC processes another command, the process outlined above is repeated.

Using DDL Statements to Create the Internal Tables

This appendix explains how to use DDL statements to create the Resource Management internal tables.

In this appendix:

- [Overview of Creating the Resource Analyzer and Resource Governor Internal Tables](#)

Overview of Creating the Resource Analyzer and Resource Governor Internal Tables

The Resource Management internal tables are created using the server Web Console. If you need or prefer to create the internal tables yourself, as required by your site, you can create them for any supported relational database.

Procedure: **How to Create the Internal Tables**

To create the Resource Management repository table:

1. Start the Web Console of the server on which Resource Management will be configured.
2. Verify that the data adapter is configured to the database where the Resource Management repository will be created. If the data adapter is not configured, you must do so now before continuing.

Note: The same Connection Name used when configuring the data adapter is required.

3. Configure Resource Management.
 - a. Click the *Resource Management* link on the sidebar.
 - b. Click *Configure* in the ribbon or right-click *Resource Management* in the navigation pane and select *Configure*.
 - c. Enter the license key from the packing slip and click *Next*.

Note: For more information, see [Configuring Resource Management](#) on page 22.

4. Enter the server name and select *Relational* from the Repository drop-down menu.
5. Select *DDL Only* from the Create Repository drop-down menu.

6. Enter the necessary information required to configure the repositories, as shown in the following image.

Configure Resource Management

Repository Settings

? Server Name

? Repository

? Data Adapter

? Create Repository

? Owner/Schema (case sensitive)

Configuration Settings

? Data Request Monitoring

? Procedure Request Monitoring

? WebFOCUS System Procedure Monitoring

? Repository load type (Before selecting Bulk load confirm that the vendor's utility has been installed)

Note: For Data Adapter and Owner/Schema, you must use the same values you used when configuring the adapter.

7. Click *Create DDL*.
8. Edit the file as needed for use at your site:
 - a. Add any additional site-specific requirements. Consult the appropriate RDBMS manuals for additional information.
 - b. Using a command shell, or another product, execute the modified DDL to create the Resource Management repository tables.

Note: The DDL statements are created without a semi colon at the end of each CREATE command. If your platform requires this, for example, Mainframe Db2 via ISPF SPUFI, the semi colon must be added. Do not add it if your environment is UDB.
9. Configure Resource Management.
 - a. Click the *Resource Management* link on the .
 - b. Click *Configure* in the ribbon or right-click *Resource Management* in the navigation pane and click *Configure*.
 - c. Enter the license key from the packing slip and click *Continue*.

Note: For more information, see [Configuring Resource Management](#) on page 22
10. Select *No* from the Create Repository drop-down menu.

11. Enter the same information that you used in step 6.
12. Click *Configure*.

Business Rule Language (BRL)

Resource Governor lets you customize rules and messages for a particular user or situation using a language called Business Rule Language (BRL).

In this appendix:

- [Using Business Rule Language \(BRL\)](#)
 - [BRL Factual Information](#)
 - [BRL Keywords](#)
 - [Pre-defined Variables](#)
 - [Internal Functions](#)
 - [Customized Rule Examples](#)
-

Using Business Rule Language (BRL)

Business Rule Language (BRL) allows IF/THEN testing on certain information available when the rules execute. Your own cancel message can be included within the rules you create by using special message variables.

BRL is a non-procedural, high-level application development language that allows you to develop sophisticated programs with less effort than conventional programming languages. BRL rules consist of three parts: the rule name, which serves as a comment or description, and is not syntactically necessary to the rule; a supporting condition (antecedent) or procedure statement; and a conclusion.

The procedure and conclusion are expressed through IF-THEN statements. When a rule is executed, Resource Governor determines whether these IF-THEN statements lead to a specified goal. Every BRL rule file, or knowledge base, must have at least one goal statement describing a conclusion that can be reached by your rules and that affects governing decisions, deciding whether the query is to be canceled or run. The final goal statement pre-established for all Resource Governor rules is *DBA Rules Concluded*, but you may also specify additional sub-goals within your customized rule file.

Resource Governor uses a process called *backward chaining*, which involves starting with the final goal and working *backward* through the sub-goals expressed within the rules to arrive at the final goal.

BRL Factual Information

BRL is capable of representing several different types of information within a single rule file or rule:

- Simple factual statements.
- Numeric data.
- String-type data.
- Attribute-value (A-V) associations.

These facts are described in the following sections. BRL facts are case sensitive and must be used exactly as shown. These sections give examples based on the following rule:

```
RULE          Check Conditions
IF           Select asterisk
AND         Day of the Week <5
AND         TableName IS XXX
AND         Current Date = "20080709"
THEN        DBA Rules Concluded
```

Reference: Simple Facts

A simple fact is a straightforward expression that requires only a true or false answer. An example of a simple fact is "Select asterisk."

Reference: Numeric Facts

Numeric data lets you compare values, create computations, and assign numeric data to variables. BRL recognizes that you are specifying a numeric data type by the presence of one of the reserved relational, numerical, or assignment operators within a supporting condition (antecedent, or IF statement).

Relational Operators	<	Less than
	>	Greater than
	<=	Less than or equal to
	>=	Greater than or equal to
	< >	Not equal
	=	Equal
Assignment Operator	:=	

When you are using numeric facts, all variables must be declared as a NUMERIC type.

An example of a numeric fact is, "Day of the Week."

Reference: String Facts

With string facts, you can pass string information to external programs, construct customized messages, write customized external files, and so on. The target variable must be pre-declared as a STRING type. An example of a string fact is, Current Date = "20080709".

Reference: Attribute-Value Facts

An attribute-value association is a statement in which an attribute is described by a particular value. In this way, it resembles a simple fact, except that a simple fact has only one value associated with the attribute, while the attribute-value association uses variables to allow for a variety of values. An example of an attribute-value fact is, "AND TableName IS XXX," since there are several different values that could be substituted for "TableName" after the "IS."

Example: Using BRL

This example of BRL is more complex than the previous one, showing actual syntax for a server for Windows. The BRL shown here allows the DBA to issue a message to end users when specific data sources are inaccessible because of maintenance or system unavailability. It also illustrates how business rules work together with Resource Governor's automatically generated rules.

For this example, the DBA would have to type the list of one or more unavailable data sources into a sequential data set (in this case, C:\temp\unavail.dat). In addition, the DBA would enter a custom message into a file named C:\temp\appname_datasourcename.dat (see the code below for an example). This BRL contains examples of how to detect full-table scans and select *s, two operations that may be expensive to run. It also illustrates how to establish user overrides and Cartesian product join detection, as well as a number of other capabilities.

```

RULE Conclude DBA Rules Main                                |Rule1
IF Database Unavailable
OR Nocancel
OR Omnipotent User Override
OR Cartesian Product Join
OR Full Table Scan
OR Selected All Columns
OR Too Many Joins in affect
OR Valid Cancel
THEN DBA Rules Concluded
!
! Check to see if the Database is available

```

```

RULE Check If Database Unavailable                          |Rule 2
IF Initialize and Allocate
AND Unavailable
AND CloseFile
AND Lineptr := 1
AND Open custom message file
AND Read and Write Messages
AND Deallocate
THEN Database Unavailable
AND Run := "F"
AND Reason := "DBDown"
!
! CloseFile

```

```

RULE Closefine                                             |Rule3
THEN CloseFile
AND ACTIVATE %CLOSE
DATA DDN 1
DATA ioresult
!
! Initialize and Allocate

```

```

RULE Initialize and Allocate
THEN Initialize and Allocate
AND ioreult := 0
AND DDN 1 := "UNAVAIL"
AND Dynam Data 1 := " DISK C:\temp\unavail.dat"
AND ACTIVATE %FILEDEF
DATA DDN 1
DATA Dynam Data 1
!
! Unavailable

```

IRule 4

```

RULE Unavailable
IF Get Table Names
AND Read List
AND ioreult = 0
THEN Unavailable
!
! Build string for FILEDEF
! For example if ibisamp/car is current Table value
! FILEDEF MSGFILE DISK C:\temp\ibisamp_car.dat
!

```

IRule5

```
RULE Open custom message file
THEN Open custom message file
AND ioreult := 0
AND DDN 2 := "MSGFILE"
AND Dynam Data := " DISK C:\temp\"
! Need to convert / to _ before FILEDEF of individual message file
AND Extra String := "/"
AND ACTIVATE %LOCATE
DATA Record
DATA Extra String
DATA Extra Number
AND End Pos := Extra Number - 1
AND Start Pos := 1
AND ACTIVATE %SUBSTR
DATA Dynam Data 2
DATA Record
DATA Start Pos
DATA End Pos
AND Start Pos := Extra Number + 1
AND End Pos := 999
AND ACTIVATE %SUBSTR
DATA Dynam Data 3
DATA Record
DATA Start Pos
DATA End Pos
AND Extra String := "_"
AND ACTIVATE %CONCAT
DATA Dynam Data 2
DATA Extra String
AND ACTIVATE %CONCAT
DATA Dynam Data 2
DATA Dynam Data 3
AND ACTIVATE %CONCAT
DATA Dynam Data
DATA Dynam Data 2
AND Extra String:= ".dat"
AND ACTIVATE %CONCAT
DATA Dynam Data
DATA Extra String
AND ACTIVATE %FILEDEF
DATA DDN 2
DATA Dynam Data
!
```

IRule6

```

RULE Read and Write Messages
ACTIVATE %GET
DATA DDN 2
DATA Record
DATA ioreult
IF message line 1
AND message line 2
AND message line 3
AND message line 4
AND message line 5
AND ioreult <> 0
THEN Read and Write Messages
ELSE FORGET message line 1
AND FORGET message line 2
AND FORGET message line 3
AND FORGET message line 4
AND FORGET message line 5
AND Lineptr := Lineptr + 1
AND LOOP
!
```

IRule7

```

RULE to build message line 1
IF Lineptr = 1
THEN message line 1
AND Message1 := Record
ELSE message line 1
!
```

IRule8

```

RULE to build message line 2
IF Lineptr = 2
THEN message line 2
AND Message2 := Record
ELSE message line 2
!
```

IRule9

```

RULE to build message line 3
IF Lineptr = 3
THEN message line 3
AND Message3 := Record
ELSE message line 3
!
```

IRule 10

```

RULE to build message line 4
IF Lineptr = 4
THEN message line 4
AND Message4 := Record
ELSE message line 4
!
```

IRule 11

```
RULE to build message line 5
IF Lineptr = 5
THEN message line 5
AND Message5 := Record
ELSE message line 5
!
```

IRule 12

```
RULE Read List
ACTIVATE %GET
DATA DDN 1
DATA Record
DATA ioreresult
ACTIVATE %TRIM
DATA Record
DATA right
IF Database Found
OR ioreresult <> 0
THEN Read List
ELSE FORGET Database Found
AND LOOP
!
```

IRule 13

```
RULE Find Unavailable Database
IF Tablename 1 = Record
OR Tablename 2 = Record
OR Tablename 3 = Record
OR Tablename 4 = Record
OR Tablename 5 = Record
THEN Database Found
!
```

IRule 14

```
RULE Get Table Names
ACTIVATE %TABLES
DATA Table IS WHAT
IF Make string 1
OR Make string 2
OR Make string 3
THEN Get Table Names
!
```

IRule 15

```
RULE Make string 1
IF Table IS ibisamp/employee
THEN Make string 1
AND Tablename 1 := "ibisamp/employee"
!
```

IRule 16

```
RULE Make string 2
IF Table IS ibisamp/car
THEN Make string 2
AND Tablename 2 := "ibisamp/car"
!
```

IRule 17


```

RULE Make string 3                                |Rule 18
IF Table IS ibisamp/jobfile
THEN Make string 3
AND Tablename 3 := "ibisamp/jobfile"
!
! Deallocate the DDNAME for unavail list and the message file

```

```

RULE Deallocate                                  |Rule 19
THEN Deallocate
AND Dynam Data := "CLEAR"
AND Dynam Data 1 := DDN 1
AND ACTIVATE %FILEDEF
DATA Dynam Data 1
DATA Dynam Data
AND Dynam Data 1 := DDN 2
AND ACTIVATE %FILEDEF
DATA Dynam Data 1
DATA Dynam Data
!

```

```

RULE Valid Cancel                                |Rule 20
IF Run <> "T"
THEN Valid Cancel
AND Message1 := "Execution cannot be allowed."
AND Message2 := "Query may potentially consume excessive resources."
!
! Disable governing between 6 PM and 8 AM

```

```

RULE Nocancel                                    |Rule 21
IF Run <> "T"
AND Hours and minutes >= "1800"
AND Hours and minutes <= "0800"
THEN Nocancel
AND Run := "T"
!

```

```

RULE To allow Omnipotent users                   |Rule 22
IF Userid = "CFO "
OR Userid = "PRESIDNT"
OR Userid = "VICEPRES"
THEN Omnipotent User Override
AND Run := "T"
!
! No WHERE conditions so a FULL table scan is imminent

```

```
RULE No Where Conditions Found IRule 23
IF Number of relations = 0
THEN Full Table Scan
AND Run := "F"
AND Reason := "FullScan"
AND Message1 := "Please add a WHERE condition to your request"
AND Message2 := "to avoid a full table scan and excessive"
AND Message3 := "resource consumption."
!
```

```
RULE Selected All Columns IRule 24
IF Select asterisk
THEN Selected All Columns
AND Run := "F"
AND Reason := "SELECT*"
AND Message1 := "Please do not choose EVERY columns in this"
AND Message2 := "particularly wide database.Instead"
AND Message3 := "choose specific column names in your query."
!
! Check to see if the user is issuing a Cartesian Product with
! the join of two databases
```

```
RULE Cartesian Product Join IRule 25
IF Number of tables = 2
AND Number of relations = 0
THEN Cartesian Product Join
AND Run := "F"
AND Reason := "CartProd"
AND Message1 := "Apply a WHERE clause to constrain this"
AND Message2 := "join. Your request is a Cartesian Product"
AND Message3 := "and will return n X m rows. "
!
! The use of many joins can use excessive resources
```

```
RULE Too Many Joins in affect IRule 26
IF Number of tables >= 4
THEN Too Many Joins in affect
AND Run := "F"
AND Reason := ">=4joins"
AND Message1 := "You are using too many joined data objects."
AND Message2 := "For efficiency, extract data to temp files"
AND Message3 := "and join the temp files to persistent databases."
!
```

```
RULE Conclude DBA Rules Fallthrough Rule
THEN DBA Rules Concluded
```

Rule 1 is the main rule and is inspected first. If any one of the OR conditions fires (that is, is true), then the rule is completed and no further OR conditions of Rule 1 are inspected. The next rule, Rule 2, must fire five additional rules in order to confirm that a database is unavailable, as well as to send a message to the end user. You can follow the rules to their conclusion by reading Rules 3 through 7. Rule 7 reads a flat file and sets up messaging output (Rules 8 through 12). Rule 13 reads the list of currently unavailable databases from a flat file. Rule 20 messages the user when the Resource Governor has determined using Resource Governor's automatically generated rules that the query exceeds the user input threshold. Rules 21 through 26 address specific query characteristics. For example, Rule 25 checks to see if the request is joining two tables without an IF/WHERE condition (a Cartesian product join), which is a potentially costly and inefficient action. Rule 19 performs clean-up operations.

BRL Keywords

You can add customized rules as required by your site. All BRL keywords must be written in uppercase.

Reference: Variable Declarations

The following table lists and describes the BRL keywords for variable declarations.

Note: Variable names with an asterisk (*) after them are available for User coding and are case sensitive.

Variable Name	Type	Usage
Current IOs Threshold	NUMERIC	Reference Only
IOs Threshold	NUMERIC	Reference Only
Threshold	NUMERIC	Reference Only
Start Pos	NUMERIC	Available for User
End Pos	NUMERIC	Available for User
message line 1	SIMPLEFACT	Reference Only
message line 2	SIMPLEFACT	Reference Only
message line 3	SIMPLEFACT	Reference Only
message line 4	SIMPLEFACT	Reference Only

Variable Name	Type	Usage
message line 5	SIMPLEFACT	Reference Only
left	STRING	Reference Only
both	STRING	Reference Only
KNB Type	STRING	Reference Only
Record*	STRING	Available for user

Reference: Rule Declarations

The following table lists and describes the BRL keywords for rule declarations.

Keyword	Description
ACTIVATE	Allows execution of external routines or pre-defined internal functions. See Internal Functions on page 215 for a list of executable ACTIVATE routines.
AND	Serves as a logical operator for joining variables or a sequencing operator when it connects a series of actions or statements. In premises containing both AND and OR, statements are evaluated from the top to the bottom of the rule, in order of appearance.
ARE	The keywords IS, ARE, and the symbol \ can be used interchangeably. (See the keyword IS.)
DATA	When used with READ, causes a variable to be assigned a value that is read from memory in the server. In conjunction with WRITE, DATA causes the value of a fact to be written to memory in the server. The purpose of DATA is to transfer the value of a variable between a rule file and the server or functions invoked through the ACTIVATE keyword. Each DATA keyword must be on its own line in the rule file. The value of a variable read in by DATA goes into the session context and is useable from then on.
END	Marks the end of a rule file and signals the termination of the text of BRL. Omitting END will result in a compiler error.

Keyword	Description
IF	Introduces the premise of a rule. There can only be one IF per rule. Multiple premises can be tested and joined by AND or OR.
INIT	Enters values into the session context when a rule file is first executed. The purpose of INIT is to assign values needed when processing begins. Variables initialized with INIT are available before any rules execute.
IS	Binds an ATTRIBUTE variable to a value. In the premise of a rule, the IS value is interpreted as a test to determine whether the attribute-value pair is a true statement. In a conclusion, the IS value assigns a value of true to the statement. The keywords IS, ARE, and the symbol \ can be used interchangeably.
NOT	Negates the true value of a variable. In the premise of a rule, a statement beginning with NOT is true if the variable following NOT is false and vice versa. In the conclusion of a rule, NOT enters the following variable into the session with the value false if the rule executes.
OR	The logical operator for connecting variables in the premise of a rule. OR makes it possible to combine rules having the same conclusion into a single rule. In premises containing both AND and OR, statements are evaluated from the top to the bottom of the rule, in order of appearance.
READ	Causes values from the following DATA keywords to be assigned to the variables listed with DATA.
RULE	Specifies the name of a BRL rule. Every rule must begin with the reserved word RULE followed by the name of the rule. Rule names do not have to be distinct. All rules must contain the keyword THEN followed by a variable that is a SIMPLEFACT or ATTRIBUTE.

Keyword	Description
THEN	Introduces the primary conclusion of a rule. Every rule must have the keyword THEN. All the BRL statements in the primary conclusion are executed when the premise of the rule is true. The first statement following THEN must be a SIMPLEFACT type of variable. Numeric, string, and procedural statements are not permitted immediately following THEN although they are allowed elsewhere in the conclusion. The primary goal that the rule engine tries to solve is DBA Rules Concluded. However, any alternate SIMPLEFACT may be used if subsequent rules test for the alternate and solve for the DBA Rules Concluded. This forces back-chaining and all the rules to be evaluated.
WHAT	Used with IS, ARE, or the symbol \ and the keyword DATA to indicate that multiple values are to be returned or sent for an ATTRIBUTE (as in DATA attribute IS WHAT).
! or ?	The comment operator marks the beginning of a comment in a BRL source file.
" "	Quotation marks designate a string literal. Literals must be enclosed in double quotation marks to permit the use of reserved BRL symbols and words and leading and trailing blanks when assigning a value to a variable, and to differentiate a literal value for a variable from a variable name.
:=	The assignment operator assigns the value on the right side of the := symbol to the variable on the left side (this operator is used with NUMERIC and STRING variables).
\	The keywords IS, ARE, and the symbol \ can be used interchangeably. (See the keyword IS.)

Keyword	Description
Arithmetic Operators	<p>Used in performing calculations involving numeric expressions.</p> <p>() - Parenthesis.</p> <p>* - Multiplication.</p> <p>/ - Division.</p> <p>+ - Addition.</p> <p>- - Subtraction.</p>
Relational Operators	<p>Relational operators are not allowed in conclusions of rules. Strings or numerics can be compared.</p> <p>= - Equal to.</p> <p>< > - Not equal to.</p> <p>< - Less than.</p> <p>> - Greater than.</p> <p><= - Less than or equal to.</p> <p>>= - Greater than or equal to.</p>

Pre-defined Variables

Variables are case and white-space sensitive. For example, TIME is a different variable than time. The following sections list all the variables provided in a rule file. A MULTI can have from 1 to 256 values.

These rules apply to all of the variables listed.

NUMERIC variables are standard double-precision floating-point data types.

SIMPLEFACT variables are always true or false.

In an assignment, a premise, or a conclusion in a rule, variables are always referenced by themselves or with the NOT operator. For example, IF *simple1* or IF NOT *simple1*

STRING variables can contain up to 80 characters.

Reference: NUMERIC Variables

The following table lists and describes the NUMERIC variables available in BRL.

Variable	Description
Reference Only	
<code>Current Time Threshold</code>	The threshold in effect for elapsed or CPU time in seconds for the current date and time.
<code>Current Rows Threshold</code>	The threshold in effect for the maximum size of the requested result set.
<code>Time Threshold</code>	The threshold estimated by the rule file for elapsed or CPU seconds.
<code>Rows Threshold</code>	The threshold estimated by the rule file for the maximum size of result set.
<code>Day of Week</code>	The day of the week indicated as 0 - 6, where 0 is Sunday and 6 is Saturday.
<code>Number of unions</code>	The number of UNION statements contained in the request.
<code>Number of union alls</code>	The number of UNION ALL statements contained in the request.
<code>Number of tables</code>	The number of tables or data sources contained in the request.
<code>Number of relations</code>	The number of relational clauses contained in the request.
<code>Number of group bys</code>	The number of GROUP BY columns contained in the request.
<code>Number of order bys</code>	The number of ORDER BY columns contained in the request.
<code>Number of functions</code>	The number of functions contained in the request.
<code>Number of columns</code>	The number of columns contained in the request or records inserted into the table SMCOLUMNS. It is primarily the number of columns selected except when an asterisk is used. In that case, the number of columns is one.

Variable	Description
<code>Record Limit</code>	Contains the RECORDLIMIT value if any was used with the TABLE syntax in the current report request being governed.
Available for User	
<code>Rule Number</code>	An arbitrary number assigned to a rule to enable the reference of what rule did or did not cancel a request.
<code>Extra Number</code>	Defined for use by business rules.
<code>Ioresult</code>	Available for any use initialized to zero.
<code>Lineptr</code>	Available for any use initialized to zero.

Reference: SIMPLEFACT Variables

The following table lists and describes the SIMPLEFACT variables available in BRL.

Variable	Description
Reference Only	
<code>Correlated query</code>	Used by the automated rules created with GKERULE, and can be used in business rules as well. Contains true when the current request is relational and correlated.
<code>Select asterisk</code>	Used by the automated rules created with GKERULE, and can be used in business rules as well. Set to true if an asterisk was used to SELECT or PRINT from a request.
<code>Bad or not condition</code>	Used with the %ORNOTERR function. Identifies whether an invalid logical combination of not/or conditions are specified in the current request WHERE clause.
<code>Distinct columns</code>	Used with the %DISTINCT function. Specifies whether the request DISTINCT parameter was used in the current query.
<code>Temporary field</code>	Contains true or false for any temporary computes or defines issued at runtime in the request.

Variable	Description
<code>Relational source</code>	Contains true for relational data requested.

Reference: STRING Variables

The following table lists and describes the STRING variables available in BRL.

Variable	Description
Reference Only	
<code>User ID</code>	The ID of the client connected to the server.
<code>Current Date</code>	The current date in the form YYYYMMDD.
<code>Hours and minutes</code>	The current time in the form HHMM.
<code>Right</code>	An extra string variable available initialized to "R".
<code>Shift</code>	Contains the current governor parameter name that is active.
<code>Request Type</code>	Contains the current request type being processed. "T" for TABLE, "S" for SQL SELECT. See more in Appendix B under table name SMQUERY, column name SMRECTYPE.
<code>Procedure</code>	Contains the name of the currently running procedure.
<code>Connection Address</code>	The IP address of the connecting process.
<code>HOLD Format</code>	The FORMAT used with the HOLD command, in TABLE syntax if used, or blank.
<code>Monitor</code>	Contains the current Monitor Preference setting.
Available for User	

Variable	Description
Reason	An eight-character string that contains a reason why the Governor decided what it did. This value is stored in the Usage Monitor table, SMGOVERN.
Message1	The first string defined that can contain a custom message. It is always returned to the server. It is used if there is a value assigned to it and overrides the standard cancellation message.
Message2	The second string defined that can contain a custom message. It is always returned to the server. It is used if there is a value assigned to it and Message1.
Message3	The third string defined that can contain a custom message. It is always returned to the server. It is used if there is a value assigned to it and the preceding message strings.
Message4	The fourth string defined that can contain a custom message. It is always returned to the server. It is used if there is a value assigned to it and the preceding message strings.
Message5	The fifth string defined that can contain a custom message. It is always returned to the server. It is used if there is a value assigned to it and the preceding message strings.
Format	A string value unused at this time. Available for any use in business rules.
Extra String	A string value unused at this time. Available for any use in business rules.
Run	The indicator for whether the decision is to run or cancel the request. It is initialized to <i>T</i> . Business rules can set it to <i>F</i> if desired, which causes the request to cancel.
DDN 1	Available for any use in business rules.
DDN 2	Available for any use in business rules.
DDN 3	Available for any use in business rules

Variable	Description
DDN 4	Available for any use in business rules
DDN 5	Available for any use in business rules
Tablename 1	Available for any use in business rules
Tablename 2	Available for any use in business rules
Tablename 3	Available for any use in business rules
Tablename 4	Available for any use in business rules
Tablename 5	Available for any use in business rules
Dynam data	Available for any use in business rules
Dynam data 1	Available for any use in business rules
Dynam data 2	Available for any use in business rules
Dynam data 3	Available for any use in business rules
Dynam data 4	Available for any use in business rules.
App Name 1	Available for use with the %APPNAME function in business rules.
App Name 2	Available for use with the %APPNAME function in business rules.
App Name 3	Available for use with the %APPNAME function in business rules.
App Name 4	Available for use with the %APPNAME function in business rules.
App Name 5	Available for use with the %APPNAME function in business rules.
Service Name	Available for use with the %SVCNAME function in business rules.

Reference: MULTI ATTRIBUTE Variables

The following table lists and describes the MULTI ATTRIBUTE variables available in BRL.

Variable	Description
Reference Only	
<code>AppName</code>	Used with the %APPNAME function. Its values are the application name used in the request.
<code>TableName</code>	Used by the automated rules created with GKERULE.
<code>Relation</code>	Used by the automated rules created with GKERULE.
<code>Table</code>	Used with the %TABLES function. Its values are the table names specified in the request.
<code>Column name</code>	Used with the %COLUMNS function. Its values are the column names used in the request.
<code>Relational</code>	Used with the %RELOPCOL function. Its values are the relational statements from the request WHERE clause with the right-hand side specified as LITERAL or FIELD. For example, <pre>WHERE (SMQUERY.SMKEY = SMFROMS.SMKEY) AND (SMQUERY.SMDATE > '19960101')</pre> would become the values: <pre>SMQUERY.SMKEY.EQ.FIELD</pre> and <pre>SMQUERY.SMDATE.GT.LITERAL</pre>
<code>Left relational column</code>	Used with the %LEFTREL function. Its values are the left-hand side of the relational statements specified in the request WHERE clause. Using the previous example, the values would be: <pre>SMQUERY.SMKEY</pre> and <pre>SMQUERY.SMDATE</pre>

Variable	Description
<code>Right relational column</code>	Used with the %RIGHTREL function. Its values are the right-hand side of the relational statements specified in the request WHERE clause when a field is specified. Using the previous example, the only value used would be: <code>SMQUERY.SMKEY</code>
<code>Group by column</code>	Used with the %GROUPBY function. Its values are the column names specified in the request GROUP BY clause of the current query.
<code>Order by column</code>	Used with the %ORDERBY function. Its values are the column names specified in the request BY clause of the current query.
<code>Function name</code>	Used with the %FUNCTION function. Its values are the function names specified in the request WHERE clause of the current query.
<code>Function column</code>	Used with the %FUNCOLS function. Its values are the function and column names specified in the request WHERE clause of the current query.
<code>Relational operator</code>	Used with the %RELOPS function. Its values are the relational operators specified in the request WHERE clause. Using the previous example, the values EQ.FIELD and GT.LITERAL would be used.
<code>Relational statement</code>	Used with the %RELATION function. Its values are the entire relational statements from the request WHERE clause. Using the previous example, the values used would be <code>SMQUERY.SMKEY.EQ.creator.SMFROMS.SMKEY</code> and <code>SMQUERY.SMDATA.GT.19960101</code>

Variable	Description
<code>Table relations</code>	Used with the %RELTABS function. Its values are the table names relation from the request WHERE clause. Using the previous example, the value used would be: <code>SMQUERY.EQ.creator.SMFROMS</code>

Internal Functions

Reference: Activate Functions

Resource Governor provides functions that are designed to assist with application development using Custom rules. These include String Functions, Date Functions, File Functions, and Current Request Functions. These functions are described below, along with variables and syntax examples. The ACTIVATE command can be executed directly from any Business rules file.

The following table lists and describes internal functions executable as ACTIVATE routines in the rules.

Function	Description
<code>%APPNAME</code>	Returns Application names.
<code>%CLOSE</code>	Closes a file defined with DYNAM or FILEDEF.
<code>%COLUMNS</code>	Returns column names.
<code>%CONCAT</code>	Concatenates two strings into a resultant string value.
<code>%COPY</code>	Makes a copy of a file.
<code>%DATE</code>	Creates a formatted date string value.
<code>%DELETE</code>	Deletes characters between starting and ending positions inclusively.
<code>%DISTINCT</code>	Returns if DISTINCT columns were requested.

Function	Description
<code>%DOW</code>	Creates a day string value from a numeric indicator.
<code>%DYNAM</code>	Uses the DYNAM command for file manipulation in MVS.
<code>%FILEDEF</code>	Uses the FILEDEF system command for reserving, deleting, or modifying files on systems other than MVS.
<code>%FTOA</code>	Converts a NUMERIC fact value to a STRING.
<code>%FUNCOLS</code>	Returns functions with column names.
<code>%FUNCTION</code>	Returns function names.
<code>%GET</code>	Reads a line from a file defined with DYNAM or FILEDEF.
<code>%GROUPBY</code>	Returns group names.
<code>&INSERT</code>	Inserts characters between starting and ending positions inclusively.
<code>%LEFTREL</code>	Returns relational columns (left side).
<code>%LENGTH</code>	Returns the length of a string value.
<code>%LOCATE</code>	Locates the existence of a string within a string value and returns the position.
<code>%LOWCASE</code>	Sets a string value to lower case.
<code>%MONTH</code>	Creates a month string value.
<code>%ORDERBY</code>	Returns ordered columns.
<code>%ORNOTERR</code>	Invalid combination of OR and NOT logic.
<code>%OVERLAY</code>	Overlays one string with another at a specified position.
<code>%PAD</code>	Pads a given string with a given character to a specified length.

Function	Description
%PUT	Writes a record to a file defined with DYNAM or FILEDEF.
%RELATION	Returns the relational statements.
%RELCNT	Returns the number of times a column name is used in a relational clause.
%RELOPCOL	Returns the relational column/operator combinations and determines whether they compare against a literal or a field.
%RELOPS	Returns relational operator AND/OR with LITERAL or FIELD.
%RELTABS	Returns relational conditions between tables.
%REQUEST	Writes the current command request being processed to a file.
%RIGHTREL	Returns relational columns (right side).
%SUBSTR	Picks out a section of a string value with starting and ending numeric positions and puts that section into a resultant string value.
%SVCNAME	Returns the service name that the current request is running under.
%TABLES	Returns table names.
%TIME	Creates a formatted time string value.
%TRIM	Trims blanks from the beginning, end, or both from a string value.
%TRUNC	Truncates a string at the desired position.
%UPCASE	Sets a string value to upper case.

Note: All function names begin with a percent (%) sign and must be entered in uppercase.

Syntax:

```
ACTIVATE function_name DATA parm 1
ACTIVATE function_name
DATA parm 1
DATA parm 2
.
.
.
DATA parm n
```

where:

function_name

Is any valid internal function name.

Parameters passed to and received from functions are identified through the DATA statement(s) as shown in the example below.

parm n

Represents the parameters being passed to *function_name* using the DATA statement.

Note: The ACTIVATE and DATA keywords must always be entered in uppercase.

Examples in the following function descriptions are based on this sample TABLE request:

```
JOIN INDEX1 IN TABLE1 TO ALL INDEX2 IN TABLE2
END
TABLE FILE TABLE1
SUM AVE.COLUMN4 MAX.COLUMN5
BY COLUMN1
BY COLUMN2
WHERE COLUMN3 GT '20101231'
END
```

Note: TABLE1 contains columns 1 and 3. TABLE2 contains columns 2, 4, and 5.

Reference: %APPNAME

The %APPNAME function will return a list of applications that the current request references.

Syntax:

```
ACTIVATE %APPNAME
DATA App Name IS WHAT
```

Example:

```

RULE Check App Names
ACTIVATE %APPNAME
DATA AppName IS WHAT
IF NOT AppName IS baseapp
THEN DBA Rules Concluded
AND RUN := "F"
AND Rule Number := 1
AND Reason := "APPNAME"
AND Message1 := "You can only run requests for Data Sources"
AND Message2 := "in BASEAPP"

```

Reference: %CLOSE

The %CLOSE function closes a file defined with DYNAM or FILEDEF.

Syntax:

```

ACTIVATE %CLOSE
DATA (input string variable)

```

Example:

```

RULE close the file
IF read record
OR write record
THEN DBA Rules Concluded
AND ACTIVATE %CLOSE
AND DATA DDN 1

```

Reference: %COLUMNS

The %COLUMNS function returns all the column names specified in the request.

Syntax:

```

ACTIVATE %COLUMNS
DATA Column name IS WHAT

```

Example:

```

RULE Get column names
ACTIVATE %COLUMNS
DATA Column name IS WHAT
THEN Column names retrieved

```

Based on the sample request, the following attribute-value pairs are true:

```

Column name IS TABLE1.COLUMN1
Column name IS TABLE2.COLUMN2

```

Reference: %CONCAT

The %CONCAT function concatenates two strings into a resultant string value.

Syntax:

```
ACTIVATE %CONCAT  
DATA (input/output string variable)  
DATA (input string variable)
```

Example:

```
RULE Concatenate strings  
ACTIVATE %CONCAT  
DATA Record  
DATA Extra String
```

Reference: %COPY

The %COPY function makes a copy of a file.

Syntax:

```
ACTIVATE %COPY  
DATA (output numeric variable)
```

Example:

```
RULE Copy file  
ACTIVATE %COPY  
DATA ioreult
```

This function will copy a file allocated as SYSUT1 to a file allocates as SYSUT2 using %DYNAM or %FILEDEF. Using the example above, if ioreult numeric variable is not a zero then the copy failed.

Reference: %DATE

The %DATE function converts month, day, and year (numeric values) to an nn/nn/nn formatted string.

Syntax:

```
ACTIVATE %DATE  
DATA (output string variable)  
DATA (input numeric variable)  
DATA (input numeric variable)  
DATA (input numeric variable)
```

Example:

```
RULE make a date
ACTIVATE %DATE
DATA Message1
DATA Message2
DATA Message3
```

Given the month, day, and year (numeric values), the syntax example updates Extra String with a string formatted mm/dd/yy, dd/mm/yy, or yy/mm/dd.

The order of the DATA statements can be changed if a different date format is required. For example, if the required format was dd/mm/yy, then day would be on the first DATA statement and month on the second.

Reference: %DELETE

The %DELETE function deletes characters between starting and ending positions inclusively.

Syntax:

```
ACTIVATE %DELETE
DATA (input/output string variable)
DATA (input string variable)
DATA (input string variable)
```

Example:

```
RULE Delete characters
ACTIVATE %DELETE
DATA Extra String
DATA Start Pos
DATA End Pos
```

The syntax example deletes all characters in Extra String between the specified Start Pos and End Pos, inclusively.

Reference: %DISTINCT

The %DISTINCT function returns true if DISTINCT is specified in the SELECT clause; otherwise, it returns false.

Syntax:

```
ACTIVATE %DISTINCT
DATA Distinct columns
```

Example:

```
RULE Get distinct
ACTIVATE %DISTINCT
DATA Distinct columns
THEN Distinct determined
```

Based on the sample request, the SIMPLEFACT Distinct columns are false. Distinct columns would be true if the sample had been:

```
PRINT DST.COLUMN1 ...
```

Reference: %DOW

The %DOW function creates a day string value from a numeric indicator. The Day of Week NUMERIC value passed into business rules is a number from 0 through 6, Sunday through Saturday.

Syntax:

```
ACTIVATE %DOW
DATA (input numeric variable)
DATA (output string variable)
```

Example:

```
RULE day name
ACTIVATE %DOW
DATA Day of Week
DATA Extra String
```

Reference: %DYNAM

The %DYNAM function is used to manipulate files in PDS deployment.

Syntax:

```
ACTIVATE %DYNAM
DATA (input string variable)
DATA (input string variable)
```

Example:

```
RULE Allocate The CNTLFILE
THEN Allocate The CNTLFILE
AND DDN 1 := "CNTLFILE"
AND Dynam Data := "FILE CNTLFILE"
AND Dynam Data 1 := "DA SYS1.MY.CNTLFILE SHR REU"
AND ACTIVATE %DYNAM
AND DATA Dynam Data
AND DATA Dynam Data 1
```

Reference: %FILEDEF

The %FILEDEF function is equivalent to using the %DYNAM command and uses the FILEDEF system command for reserving, deleting, or modifying files on systems other than PDS deployment.

Syntax:

```
ACTIVATE %FILEDEF
DATA (input string variable)
DATA (input string variable)
```

Example:

```
RULE Allocate The CNTLFILE
THEN Allocate The CNTLFILE
AND DDN 1 := "CNTLFILE"
AND Dynam Data 1 := "DISK C:\\CNTLFILE.TXT"
AND ACTIVATE %FILEDEF
AND DATA DDN 1
AND DATA Dynam Data 1
```

Reference: %FTOA

The %FTOA function converts a NUMERIC fact value to a STRING.

Syntax:

```
ACTIVATE %FTOA
DATA (output string variable)
DATA (input string variable)           !Format String
DATA (input numeric variable)
```

Example:

```
RULE convert to alpha
Extra String 1 := "4.2"
Extra Number := 3.1294
ACTIVATE %FTOA
DATA Extra String
DATA Extra String 1
DATA Extra Number
```

Format String is in the form *n.d*, where *n* represents the total number of characters (including decimal point and sign) of the output string, and *d* represents the number of decimal places. A value of *n=0* indicates that the number of places to the left of the decimal point will be as large as necessary to hold the entire number, but no larger than needed.

Reference: %FUNCOLS

The %FUNCOLS function returns all the functions with column names specified in the request.

Syntax:

```
ACTIVATE %FUNCOLS
DATA Function column IS WHAT
```

Example:

```
RULE Get functions with columns
ACTIVATE %FUNCOLS
DATA Function column IS WHAT
THEN Functions with columns retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Function column IS AVG.TABLE2.COLUMN4
Function column IS MAX.TABLE2.COLUMN5
```

Reference: %FUNCTION

The %FUNCTION function returns all the function names specified in the request.

Syntax:

```
ACTIVATE %FUNCTION
DATA Function name IS WHAT
```


Example:

```
RULE Get functions
ACTIVATE %FUNCTION
DATA Function name IS WHAT
THEN Functions retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Function name IS AVG
Function name IS MAX
```

Reference: %GET

The %GET function reads a line from a file defined with DYNAM or FILEDEF.

Syntax:

```
ACTIVATE %GET
DATA (input string variable)
DATA (output string variable)
DATA (output numeric variable)
```

Example:

```
RULE read a record
ACTIVATE %GET
DATA DDN 1
DATA Record
DATA ioreresult
IF ioreresult = 0
THEN read record
```

Reference: %GROUPBY

The %GROUPBY function returns all the column names specified in the GROUP BY clause.

Syntax:

```
ACTIVATE %GROUPBY
DATA Group by column IS WHAT
```

Example:

```
RULE Get group bys
ACTIVATE %GROUPBY
DATA Group by column IS WHAT
THEN Group bys retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Group by column IS TABLE1.COLUMN1
Group by column IS TABLE2.COLUMN2
```

Reference: %INSERT

The %INSERT function inserts characters between starting and ending positions inclusively.

Syntax:

```
ACTIVATE %INSERT
DATA (input/output string variable)
DATA (input string variable)
DATA (input numeric variable)
```

Example:

```
RULE Insert characters
ACTIVATE %INSERT
DATA Extra String
DATA Extra String 1
DATA Start Pos
```

The syntax example inserts characters in Extra String1 into Extra String at the position specified in Start Pos.

Reference: %LEFTREL

The %LEFTREL function returns the left-hand side of the relational condition as specified in the WHERE clause.

Syntax:

```
ACTIVATE %LEFTREL
DATA Left relational column IS WHAT
```

Example:

```
RULE Get left relational column
ACTIVATE %LEFTREL
DATA Left relational column IS WHAT
THEN Left relational column retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Left relational column IS TABLE1.INDEX1
Left relational column IS TABLE1.COLUMN3
```

Reference: %LENGTH

The %LENGTH function returns the length of a string value.

Syntax:

```
ACTIVATE %LENGTH
DATA (output numeric variable)
DATA (input string variable)
```

Example:

```
RULE find string length
ACTIVATE %LENGTH
DATA Extra Number
DATA Record
```

The output string variable contains the length of the input string variable.

Reference: %LOCATE

The %LOCATE function locates the existence of a string within a string value and returns the position.

Syntax:

```
ACTIVATE %LOCATE
DATA (input string variable)
DATA (input string variable)
DATA (output numeric variable)
```

Example:

```
RULE locate a string
ACTIVATE %LOCATE
Tablename 1 := "Mercedes Benz"
Tablename 2 := "Benz"
DATA Tablename 1
DATA Tablename 2
DATA Extra Number
```

The numeric value of the third DATA statement contains the index value of the second DATA statement variable in the first DATA statement variable. If the value of the second DATA statement is not found, a zero will be returned in the third DATA statement.

Reference: %LOWCASE

The %LOWCASE function sets a string value to lower case.

Syntax:

```
ACTIVATE %LOWCASE
DATA (input/output string variable)
```

Example:

```
RULE lower case
Extra String := "case change"
ACTIVATE %LOWCASE
DATA Extra String
```

Reference: %MONTH

The %MONTH function creates a month string value from a numeric indicator 0 through 12, January through December.

Syntax:

```
ACTIVATE %MONTH
DATA (input numeric value)
DATA (output string value)
```

Example:

```
RULE month name
ACTIVATE %MONTH
DATA Extra Number
DATA Extra String
```

Reference: %ORDERBY

The %ORDERBY function returns all the column names specified in the ORDER BY clause.

Syntax:

```
ACTIVATE %ORDERBY
DATA Order by column IS WHAT
```

Example:

```
RULE Get group bys
ACTIVATE %ORDERBY
DATA Order by column IS WHAT
THEN Order bys retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Order by column IS TABLE1.COLUMN1
Order by column IS TABLE2.COLUMN2
```

Reference: %ORNOTERR

The %ORNOTERR function returns true if a faulty OR/NOT condition is specified in a WHERE clause.

Syntax:

```
ACTIVATE %ORNOTERR
DATA Bad or not condition
```

Example:

```
RULE Determine faulty relational condition
ACTIVATE %ORNOTERR
DATA Bad or not condition
THEN Faulty relational determined
```

Based on the sample request, the SIMPLEFACT Bad or not condition are false. If the sample had been:

```
TABLE FILE TABLE
PRINT COLUMN1
WHERE COLUMN9 NE 'RED' OR 'BLUE'
END
```

Then, when COLUMN9 is RED, it is not BLUE, so the OR makes the condition true; and when COLUMN9 is BLUE, it is not RED, so the OR still makes the condition true; and when COLUMN9 is neither RED nor BLUE, the condition is still true. The effect is the same as coding PRINT COLUMN1 without a WHERE clause.

In this case, the Bad or not condition is true.

Reference: %OVERLAY

The %OVERLAY function overlays one string with another at a specified position.

Syntax:

```
ACTIVATE %OVERLAY
DATA (input/output string variable)
DATA (input string variable)
DATA (input numeric variable)
```

Example:

```
RULE Overlay a string
ACTIVATE %OVERLAY
DATA Extra String
DATA Extra String 1
DATA Start Pos
```

The syntax example overlays Extra String with Extra String 1 at position Start Pos for the length of Extra String 1. If the starting position is relative to 1, for example, the first position in the string that can be overlaid is position = 1.

Reference: %PAD

The %PAD function pads a given string with a given character to a specified length.

Syntax:

```
ACTIVATE %PAD
DATA (input/output string variable)
DATA (input string variable)
DATA (input numeric variable)
DATA (defined string variable)
```

Example:

```
RULE Pad string
Extra String 1 := "x"
Extra Number = 3
ACTIVATE %PAD
DATA Extra String
DATA Extra String 1
DATA Extra Number
DATA right
```

The syntax example pads Extra String with the character in Extra String 1 for a total of Extra Number characters.

The fourth DATA statement must be one of the system defined string variables:

```
left      will trim characters from the left side of the string.
right     will trim characters from the right side of the string.
both      will trim characters from both the left and right sides of the
string.
```

Reference: %PUT

The %PUT function writes a record to a file defined with DYNAM or FILEDEF.

Syntax:

```
ACTIVATE %PUT
DATA (input string variable)
DATA (input string variable)
DATA (output numeric variable)
```

Example:

```
RULE Write a record
ACTIVATE %PUT
DATA DDN 1
DATA Record
DATA ioreult
IF ioreult = 0
THEN write record
```

Reference: %RELATION

The %RELATION function returns all the relational statements specified in the WHERE clause.

Syntax:

```
ACTIVATE %RELATION
DATA Relational statement IS WHAT
```

Example:

```
RULE Get relational statements
ACTIVATE %RELATION
DATA Relational statement IS WHAT
THEN Relational statements retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Relational statement IS TABLE1.INDEX1.EQ.TABLE2.INDEX2
Relational statement IS TABLE1.COLUMN3.GT.19951231
```

Reference: %RELCNT

The %RELCNT function returns the number of times a column name is used in a relational clause.

Syntax:

```
ACTIVATE %RELCNT
DATA (input string variable)
DATA (output numeric variable)
```

Example:

```
RULE Count column use in relational tests
Extra String 1 := "COLUMN3"
ACTIVATE %RELCNT
DATA Extra String 1
DATA Extra Number
```

Based on the sample request, Extra Number will equal 1.

Reference: %RELOPCOL

The %RELOPCOL function returns all the relational column/operator combinations, and determines whether they compare against a literal or a field, as specified in the WHERE clause.

Syntax:

```
ACTIVATE %RELOPCOL  
DATA Relational IS WHAT
```

Example:

```
RULE Get relational column and operator  
ACTIVATE %RELOPCOL  
DATA Relational IS WHAT  
THEN Relational column and operator retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Relational IS TABLE1.INDEX1.EQ.FIELD  
Relational IS TABLE1.COLUMN3.GT.LITERAL
```

Reference: %RELOPS

The %RELOPS function returns all the relational operators and determines whether they compare against a literal or a field as specified in the WHERE clause.

Syntax:

```
ACTIVATE %RELOPS  
DATA Relational operator IS WHAT
```

Example:

```
RULE Get relational operator  
ACTIVATE %RELOPS  
DATA Relational operator IS WHAT  
THEN Relational operator retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Relational operator IS EQ.FIELD  
Relational operator IS GT.LITERAL
```


Reference: %RELTABS

The %RELTABS function returns all the relational conditions between tables as specified in the WHERE clause.

Syntax:

```
ACTIVATE %RELTABS
DATA Table relations ARE WHAT
```

Example:

```
RULE Get relational table conditions
ACTIVATE %RELTABS
DATA Table relations ARE WHAT
THEN Relational table conditions retrieved
```

Based on the sample request, the following attribute-value pairs are true:

```
Relational statement IS TABLE1.EQ.TABLE2
```

Reference: %REQUEST

The %REQUEST function writes the current command request being processed to a file.

Syntax:

```
ACTIVATE %REQUEST
DATA (input string variable)
```

Example:

```
RULE Put request in a file
ACTIVATE %REQUEST
DATA DDN 1 !DDN 1 is a file that has been previously defined and
allocated
```

Reference: %RIGHTREL

The %RIGHTREL function returns the right-hand side of the relational condition when it references a field as specified in the WHERE clause.

Syntax:

```
ACTIVATE %RIGHTREL
DATA Right relational column IS WHAT
```

Example:

```
RULE Where FieldName
ACTIVATE %RIGHTREL
DATA Right relational column IS WHAT
IF NOT Right relational column IS CREATOR.TABLE2.INDEX2
THEN DBA Rules Concluded
AND Rule Number := 444
AND Run := "F"
AND Reason := "WhereCo"
AND Message1 := "Must use WHERE with CREATOR.TABLE2.INDEX2 as right column
name"
```

Based on the sample request, the following attribute-value pairs are true:

```
Right relational column IS TABLE2.INDEX2
```

Reference: %SUBSTR

The %SUBSTR function picks out a section of a string value defined by starting and ending numeric positions and puts that section into a resultant string value.

Syntax:

```
ACTIVATE %SUBSTR
DATA (input string variable)
DATA (Output string variable)
DATA (Start position)
DATA (End position)
```

Reference: %SVCNAME

The %SVCNAME function will return the Service name that the current request is running under.

Syntax:

```
ACTIVATE %SVCNAME
DATA Service Name
```

Example:

```

RULE Check Service name
ACTIVATE %SVCNAME
DATA Service Name
IF Service Name <> "DEFAULT"
THEN DBA Rules Concluded
AND Reason := "SVCNAME"
AND Rule Number := 1
AND RUN := "F"
AND Message1 := "You can only run requests in the"
AND Message2 := "DEFAULT service"

```

Reference: %TABLES

The %TABLES function returns all the table names specified in the FROM clause.

Syntax:

```

ACTIVATE %TABLES
DATA Table IS WHAT

```

Example:

```

RULE Get table names
ACTIVATE %TABLES
DATA Table IS WHAT
IF Table IS CREATOR.TABLE1
OR Table IS CREATOR.TABLE2
THEN table name used

```

Based on the sample request, the following attribute-value pairs are true:

```

Table IS TABLE1
Table IS TABLE2

```

Reference: %TIME

The %TIME function creates a formatted time string value of xx:yy, where xx or yy can be any character values chosen. For example, Message1 and Message2 may have been initialized using %SUBSTR from Day of Week.

Syntax:

```

ACTIVATE %TIME
DATA (output string variable)
DATA (input string variable)
DATA (input string variable)

```

Example:

```
RULE Make time
ACTIVATE %TIME
DATA Extra String
DATA Message1
DATA Message2
```

Reference: %TRIM

The %TRIM function trims blanks from the beginning, end, or both of a string value.

Syntax:

```
ACTIVATE %TRIM
DATA (input/output string variable)
DATA (defined string variable)
```

Example:

```
RULE
Trim blanks
ACTIVATE %TRIM
DATA Extra String
DATA right
```

The syntax example trims blanks from Extra String based on the second DATA statement value.

The second DATA statement must be one of the system defined string variables:

```
left      will trim characters from the left side of the string.
right     will trim characters from the right side of the string.
both      will trim characters from both the left and right sides of the
string.
```

Reference: %TRUNC

The %TRUNC function truncates a string at the desired position.

Syntax:

```
ACTIVATE %TRUNC
DATA (input/output string variable)
DATA (input numeric variable)
```

Example:

```

RULE truncate
Extra String := Current Date      ! Make Extra String the year
Lineptr := 5                      ! Ends the current date in the fifth
position
ACTIVATE %TRUNC
DATA Extra String
DATA Lineptr

```

Reference: %UPCASE

The %UPCASE function sets a string value to upper case.

Syntax:

```

ACTIVATE %UPCASE
DATA (input/output string variable)

```

Example:

```

RULE Upper case
Extra String := "case change"
ACTIVATE %UPCASE
DATA Extra String

```

Customized Rule Examples

The following examples suggest more uses for business rules. The goal of every custom rule group is DBA Rules Concluded. Alternate goal variables may be used as long as they are tested in subsequent rules and DBA Rules Concluded is ultimately satisfied.

Example: Initialize All of the ATTRIBUTE Variables

This example illustrates the custom rule that would make additional information available to the rule file. All of the functions called with the ACTIVATE keyword will initialize the ATTRIBUTE variables listed with the DATA statements. This allows you to write additional custom rules.

```

! Process internal all attribute/value functions
RULE Get information about the query
ACTIVATE %TABLES
DATA TableName IS WHAT
ACTIVATE %COLUMNS
DATA Column name IS WHAT
ACTIVATE %DISTINCT
DATA Distinct columns
ACTIVATE %LEFTREL
DATA Left relational column IS WHAT
ACTIVATE %RIGHTREL
DATA Right relational column IS WHAT
ACTIVATE %RELOPCOL
DATA Relational IS WHAT
ACTIVATE %GROUPBY
DATA Group by column IS WHAT
ACTIVATE %ORDERBY
DATA Order by column IS WHAT
ACTIVATE %FUNCTION
DATA Function name IS WHAT
ACTIVATE %FUNCOLS
DATA Function column IS WHAT
ACTIVATE %RELOPS
DATA Relational operator IS WHAT
ACTIVATE %RELATION
DATA Relational statement IS WHAT
ACTIVATE %RELTABS
DATA Table relations ARE WHAT
ACTIVATE %ORNOTERR
DATA Bad or not condition
THEN All functions processed

```

Example: Checking for Table Names Used

This rule example illustrates how to check for one or more table names used in the original SQL or TABLE request. For example, this rule file might have been built for TABLE2, and you know you never want TABLE1 used with TABLE2.

```

RULE Check table name
IF All functions processed
AND TableName IS TABLE2      ! This rule file is for TABLE2
AND TableName IS TABLE1      ! TABLE1 cannot be used with TABLE2
THEN DBA Rules Concluded
AND Run := "F"
AND Reason := "TABLE"
AND Rule Number := 9999
AND Message1 := "You cannot use TABLE1 and TABLE2 in the same request."

```

Example: Making Sure a Request Runs If an Equality Condition Exists

This rule example assures that if an equality condition exists on an indexed column in a WHERE clause, the request runs.

```

RULE Check Relation
IF All functions processed
AND Relational IS TABLE1.COLUMN1.EQ.LITERAL
THEN DBA Rules Concluded
AND Run := "T"
AND Reason := "INDEX"

```

Example: Canceling Requests That Use SELECT * and Have No WHERE Clauses

This rule example cancels all requests that use a SELECT * and lack WHERE clauses. It also sets the Rule Number and Reason, which are saved in the SMGOVEND RG repository table and available for reporting. The Rule Number and Reason also display to the user in a cancel message. Message1 through Message5 may contain message information that is displayed with an advise or cancel message to the user and the edaprint.log file.

```

RULE Check for Asterisk
IF Select asterisk
AND Number of Relations = 0
THEN DBA Rules Concluded
AND Run := "F"
AND Reason := "SELECT*"
AND Rule Number := 9998
AND Message1 := "Please add a WHERE clause to your SELECT statement."
AND Message2 := "You selected every row in the table."
AND Message3 := "5 Message variables are available."

```

Example: Making Sure No Requests Are Canceled

The following rule example ensures no requests are canceled.

```

RULE Do not cancel any queries
THEN DBA Rules Concluded
AND Run := "T"

```

Example: Using a Wild Card for Column Names

These rules catch any use of HEIGHT, WEIGHT, COUNTRY or CAR from the CAR file and allow the request to run. If they are not used, the process will drop down to the second rule and cancel.

```

RULE Check for column group
IF All functions processed
AND Column name IS $IGHT
OR Column name IS C$
THEN DBA Rules Concluded
AND Run := "T"
RULE fall through
THEN DBA Rules Concluded
AND Run := "F"

```

Example: Exclude Ad Hoc Requests

This rule checks if the procedure name is blank, which indicates an ad hoc request.

```
RULE No adhoc
IF Procedure = " "
THEN DBA Rules Concluded
AND Run := "F"
```

Example: No Temporary Computes or Defines Allowed on Relational Data

This rule checks for any temporary columns. They are defined or computed at run time, not in the Master File definition. This also illustrates how to check for a relational data source. You do not need to use an ACTIVATE function to get the value of the temporary field or relational source because they are facts that are passed to the governor when the rules are run.

```
RULE check temporary fields
IF Temporary field
AND Relational source
THEN DBA Rules Concluded
AND Rule Number := 1
AND Run := "F" AND
Reason := "temp"
AND Message1 := "Cannot use a a temporary field when data source is
relational"
```

Example: Introducing a Driver Rule

This example has two independent rules, each of which could return true. The first one that does will conclude the rule file. If the first rule does not return true, then the governor will proceed to the next rule. Multiple rules within a rule file can be linked together with a driver rule using AND/OR so that they are applied together.


```

!
! Rule Wizard Generated Template
!
RULE Select Asterisk
IF Select asterisk
THEN DBA Rules Concluded
AND Rule Number := 1
AND Run := "F"
AND Reason := "Select *"
AND Message1 := "No select asterisk."
!
RULE Allow Joins
IF Number of tables > 3
THEN DBA Rules Concluded
AND Rule Number := 2
AND Run := "F"
AND Reason := "Joins"
AND Message1 := "Only 3 joins allowed."
!

```

To combine the rules so that both have to be true for a request to be governed, a driver rule is needed that calls the existing two rules and specifies as its conclusion the same conclusion as before. The driver rule is:

```

RULE Stop request if Select Asterisk with more than 3 joins
IF Have Select Asterisk
AND Too Many Joins
THEN DBA Rules Concluded
AND Rule Number := 1
AND Run := "F"
AND Reason := "BUSINESS"
AND Message1 := "No select asterisk allowed with more than 3 joins."
ELSE DBA Rules Concluded
AND Rule Number := 1
AND Run := "T"

```

The two existing rules need to be changed as follows:

```

RULE Check for Select Asterisk
IF Select asterisk
THEN Have Select Asterisk
!
RULE Check Join number
IF Number of tables > 3
THEN Too Many Joins
!

```

Notice that the conclusion for the existing rules has changed, but the conditions for each remain the same.

Example: Creating a Rule File to Follow Specific Rules

The following example will create a rule file that will apply the following rules to a query:

- Disallow any request against synonyms that start with ca characters (from any application directory).
- Only during the morning hours of 9 A.M. to noon.
- Any user ID starting with BOSS, regardless of case, is always an exception to these rules.

The complete rule file to accomplish this is:

```

RULE Global Rule One
IF User is the Boss
OR Disallow Request at this time
THEN DBA Rules Concluded
!
RULE Check User ID
Start Pos := 1
End Pos := 4
ACTIVATE %SUBSTR
DATA Extra String
DATA Userid
DATA Start Pos
DATA End Pos
ACTIVATE %UPCASE
DATA Extra String
IF Extra String = "BOSS"
THEN User is the Boss
AND Run := "T"
!
RULE Check TABLE name and time of day
IF Disallow TABLE Name
AND This is Morning
THEN Disallow Request at this time
AND Run := "F"
AND Reason := "GlobalR1"
AND Message1 := "Access to any data source starting with ca from ANY "
AND Message2 := "application directory from 9AM to NOON is restricted"
ELSE Run := "T"
!
RULE Check TABLE Name
ACTIVATE %TABLES
DATA Table IS WHAT
IF Table IS $/ca$
THEN Disallow TABLE Name
!
RULE Check time of day
IF Hours and minutes >= "0900"
AND Hours and minutes <= "1200"
THEN This is Morning
!
RULE fall through
THEN DBA Rules Concluded

```

Note:

- Since the BOSS user IDs can run at any time, it is logical to test for this first. The rule uses two built-in functions:
 - %SUBSTR to get the first four characters of the supplied string variable *Userid*.
 - %UPCASE the output of %SUBSTR.
- The next rule, which is OR with the first, acts like a driver as it combines two rules using AND.
 - Does the current (any) table name active in the request start with *ca*?
 - If so, is the time between 9 A.M. and noon?
 - If the answer is Yes to both, then this OR rule will disallow the request at this time and the conclusion that is coded in this rule will be used (AND Run := "F").

Index

[%APPNAME function 218](#)
[%CLOSE function 219](#)
[%COLUMNS function 219](#)
[%CONCAT function 220](#)
[%COPY function 220](#)
[%DATE function 220](#)
[%DELETE function 221](#)
[%DISTINCT function 221](#)
[%DOW function 222](#)
[%DYNAM function 222](#)
[%FILEDEF function 223](#)
[%FTOA function 223](#)
[%FUNCOLS function 224](#)
[%FUNCTION function 224](#)
[%GET function 225](#)
[%GROUPBY function 225](#)
[%INSERT function 226](#)
[%LEFTREL function 226](#)
[%LENGTH function 226](#)
[%LOCATE function 227](#)
[%LOWCASE function 227](#)
[%MONTH function 228](#)
[%ORDERBY function 228](#)
[%ORNOTERR function 229](#)
[%OVERLAY function 229](#)
[%PAD function 230](#)
[%PUT function 230](#)
[%RELATION function 231](#)
[%RELCNT function 231](#)

[%RELOPCOL function 232](#)
[%RELOPS function 232](#)
[%RELTABS function 233](#)
[%REQUEST function 233](#)
[%RIGHTREL function 233](#)
[%SUBSTR function 234](#)
[%SVCNAME function 234](#)
[%TABLES function 235](#)
[%TIME function 235](#)
[%TRIM function 236](#)
[%TRUNC function 236](#)
[%UPCASE function 237](#)

A

administration
 [views 113, 115](#)
administrator [19](#)

B

backward chaining [193](#)
BRL (Business Rule Language) [18, 19, 103, 104, 193, 194, 196](#)
 [attribute-value associations 195](#)
 [backward chaining 193](#)
 [customizing rules with 193](#)
 [keywords 203](#)
 [numeric data 195](#)
 [simple factual statements 194](#)

BRL (Business Rule Language) [18](#), [19](#), [103](#), [104](#),
[193](#), [194](#), [196](#)

 string-type data [195](#)

browsers help [155](#)

Build Rules view [18](#)

Business Rule Language (BRL) [193](#), [194](#), [196](#)

 attribute-value associations [195](#)

 backward chaining [193](#)

 customizing rules with [193](#)

 keywords [203](#)

 numeric data [195](#)

 simple factual statements [194](#)

 string-type data [195](#)

C

calendar [19](#)

canceling requests [19](#)

configuring resource management [22](#)

 using a FOCUS repository [22](#)

 using a relational repository [22](#)

configuring the repository [22](#)

context-sensitive help [155](#)

Custom Rule Builder [15](#), [18](#), [103](#), [104](#)

D

data access [14](#)

data source options [32](#)

data sources

 monitoring [32](#)

 under control [19](#)

data warehouses [14](#)

date ranges

 activating shifts and thresholds [113](#)

dates for reports [19](#)

days of week [113](#)

drilling down on reports [142](#)

E

EDA [14](#)

enabling and disabling monitoring [32](#)

Excel [142](#)

F

field descriptions [157](#)

formats [142](#)

frequency of use [143](#)

G

governing [15](#)

 building rules [18](#)

 reports [19](#), [150](#)

graphical reports [142](#)

H

help [141](#), [153](#), [155](#)

I

IF-THEN statements [193](#)

internal functions [215](#)

- L**
- license key [22](#)
- M**
- Main menu [19](#)
 - migrating the repository [88](#)
 - modes of analysis
 - frequency of use [143](#)
 - monitor preference [32](#)
 - resetting [32](#)
 - monitored samples summary reports [19](#)
 - monitoring [14](#), [15](#), [19](#)
 - data sources [32](#)
 - enabling and disabling [32](#)
 - relational sources [32](#)
- N**
- New Rule Parameter dialog box [113](#)
 - NUMERIC variables [207](#)
- O**
- online help [141](#), [155](#)
 - operations [15](#)
 - options [138](#)
- P**
- parameters [19](#)
 - creating shifts [113](#)
 - disabling shifts [115](#)
 - parameters [19](#)
 - running by date range [113](#)
 - PDF [142](#)
 - procedures reports
 - frequency of use [143](#)
- R**
- ranges [113](#)
 - re-sorting reports [142](#)
 - Reformat As icon [142](#)
 - relational sources [32](#)
 - monitoring [32](#)
 - reports [19](#), [138](#), [147](#)
 - administration [19](#)
 - drilling down [142](#)
 - formats [142](#)
 - frequency of use [143](#)
 - governing [19](#), [150](#)
 - graphical [142](#)
 - help [141](#)
 - Knowledgebase Properties [149](#)
 - Monitored Data Sources [147](#)
 - Monitored Procedures [148](#)
 - options [138](#)
 - re-sorting [142](#)
 - Resource Thresholds [148](#)
 - Threshold Advisor [150](#)
 - repository [22](#)
 - configuring [22](#)
 - migrating [88](#)

- requests [19](#)
 - frequency of use [143](#)
- resetting monitor preference [32](#)
- Resource Governor [14](#)
 - administration [19](#)
 - Main menu [19](#)
 - operations [15](#)
 - reports [147](#)
 - rules [18](#)
- resource management [22, 88](#)
 - configuring [22](#)
 - enabling [22](#)
 - removing [32](#)
 - repository [22](#)
 - repository maintenance [32](#)
- resource usage summary reports [19](#)
- ribbon [21](#)
 - organization [21](#)
- rule files [207](#)
 - customizing [237](#)
 - variables in [207](#)
- Rule Parameters view [113, 115](#)
- Rule wizard [18](#)
- rules [15, 18](#)
 - custom [18, 103, 104](#)
 - disabling shifts [115](#)

S

- selection criteria [19](#)

- shifts [113](#)
 - disabling [115](#)
- SIMPLEFACT variables [207](#)
- SMCOLUMNS table [157](#)
- SMDSNS table [157](#)
- SMFROMS table [157](#)
- SMFUNCTIONS table [157](#)
- SMGOVERN table [157](#)
- SMGROUPBYS table [157](#)
- SMORDERBYS table [157](#)
- SMQUERY table [157](#)
- SMRELATIONS table [157](#)
- SMREQUESTS table [157](#)
- sorting [19](#)
- state column [115](#)
- STRING variables [207](#)

T

- tables [157](#)
 - SMCOLUMNS [157](#)
 - SMDSNS [157](#)
 - SMFROMS [157](#)
 - SMFUNCTIONS [157](#)
 - SMGOVERN [157](#)
 - SMGROUPBYS [157](#)
 - SMORDERBYS [157](#)
 - SMQUERY [157](#)
 - SMRELATIONS [157](#)
 - SMREQUESTS [157](#)
 - usage monitoring [157](#)

tabular reports [142](#)

Threshold Selection Adviser reports [19](#)

thresholds [15](#), [112](#), [113](#)

time ranges [113](#)

U

usage monitoring [14](#), [15](#)

reports [19](#)

V

variables [207](#)

view rules report [19](#)



Feedback

Customer success is our top priority. Connect with us today!

Information Builders Technical Content Management team is comprised of many talented individuals who work together to design and deliver quality technical documentation products. Your feedback supports our ongoing efforts!

You can also preview new innovations to get an early look at new content products and services. Your participation helps us create great experiences for every customer.

To send us feedback or make a connection, contact Sarah Buccellato, Technical Editor, Technical Content Management at Sarah_Buccellato@ibi.com.

To request permission to repurpose copyrighted material, please contact Frances Gambino, Vice President, Technical Content Management at Frances_Gambino@ibi.com.



WebFOCUS

Resource Governor Administrator's and User's Manual
Release 7709

DN3501992.0219

Information Builders, Inc.
Two Penn Plaza
New York, NY 10121-2898