

TIBCO WebFOCUS®

Reporting Server Administration

Release 9.0.0

December 2022

DN4501056.1222



Contents

1. TIBCO WebFOCUS® Reporting Server Introduction	17
Using the Reporting Server Browser Interface	17
Reporting Server Browser Interface Ribbon.....	20
Reporting Server Browser Interface Hover Menu.....	21
Menu Options	22
User Menu Options.....	22
Configuration Overview	26
TIBCO WebFOCUS Reporting Server Profiles	27
Profile Level.....	28
Profile Commands	32
Profile Command Formats.....	33
2. TIBCO WebFOCUS Reporting Server Security	35
Security Overview	35
Authentication.....	36
Security Providers.....	36
Privileges and Other Authorizations.....	39
Encryption.....	43
Configuring Authentication	43
Configuring a New Security Provider.....	43
Configuring OPSYS Authentication.....	46
Preventing Unsecured Server Starts After Upgrades.....	47
Using Password Phrases for Authentication on z/OS.....	47
Configuring Linux and AIX Pluggable Authentication Modules (PAMs).....	49
Using Kerberos for Single Sign On.....	49
Configuring PTH Authentication.....	51
Disabling a PTH Security Provider Account.....	55
Configuring DBMS Authentication.....	56
Configuring LDAP Authentication.....	56
Configuring Custom Authentication.....	56
Encoding Passwords for a Custom Security Provider.....	57
Authenticating Users Across Multiple Security Providers.....	57
Enabling Trusted Connections.....	58

Identifying a Default Provider for Trusted Groups.....	58
Setting an Anonymous User ID.....	59
Configuring User Password Settings.....	60
Sign-in Password Expiration Warning.....	61
Configuring Privileges and Other Authorizations	62
Configuring DBMS Authorization.....	63
Configuring DBA Security.....	64
Calculating Privileges for Any Registered or Unregistered User or Group.....	64
Permissions for Server Application Files and Directories Using a Non-OPSYS Security Provider.....	64
File Permissions for an OPSYS Security Provider.....	74
Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions.....	76
Transferring File Permissions With the GRANT Privilege.....	82
Configuring Groups.....	84
Configuring Roles.....	85
Registering Users and Groups in a Role.....	88
Configuring Encryption	97
Encrypting Passwords Stored in Configuration Files.....	97
User-Defined Password Encryption and Decryption.....	97
Encrypting Data at Rest.....	99
Configuring Secure Socket Layer (SSL) Encryption for the TCP/HTTP Listener.....	100
Support for OpenSSL Protocols and Ciphers, Microsoft Algorithm Identifiers, and DH and ECDH Key Exchange.....	102
Configuring Data Encryption for a Remote TIBCO WebFOCUS Reporting Server.....	103
Encoding HTML Tags in Data.....	106
Access to Connection Information in TIBCO WebFOCUS Procedures	106
Configuring Multi-Tenant Deployment	106
Configuring Security Templates for Multi-Tenant Environments.....	107
Configuring Multi-Tenant Application Folders.....	114
Configuring Multi-Tenant Groups and Users.....	114
Configuring Multi-Tenant DBMS Connections.....	115
Configuring Security Providers for Multi-Tenant Deployment.....	115

Naming Conventions for Multi-Tenant Deployment.	115
Configuring Application File Privileges for Multi-Tenant Deployment.	115
Effect of Multi-Tenant Deployment on File Upload and HOLD Files.	116
Shared Metadata Under Multi-Tenant Deployment.	116
3. Managing Applications	117
What Is an Application?	117
Generating Samples and Tutorials.	120
Procedures and Metadata on the Application Tree	121
Managing Applications and Paths	132
Creating and Mapping Applications.	132
Indexing Applications For Global Searches.	137
Indexing Options.	137
Search Index Feature.	137
Indexing Options.	138
Using an SQL Database to Store Application Contents.	143
Linking to Your TIBCO WebFOCUS Client Repository.	146
Linking to a Google Drive Repository.	148
Nested Application Directories.	150
Home Application Directories for Users.	153
Sharing Files and Applications.	156
Configuring the Application Path.	160
Filtering the Application Tree.	163
Searching for Files.	166
Sorting the Application Tree.	167
Selecting Application Files.	167
Indexing Applications For Global Searches	168
Indexing Options.	137
Search Index Feature.	169
Indexing Options.	138
Application Commands Overview	174
Search Path Management Commands	177
APP PATH.	178

APP PREPENDPATH.....	178
APP APPENDPATH.....	179
APP MAP.....	180
APP SET METALLOCATION_SAME.....	182
APP ? METALLOCATION_SAME.....	183
APP SHOWPATH.....	183
Application and File Management Commands	183
APP CREATE.....	184
APP COPY.....	186
APP COPYF[ILE].....	187
APP MOVE.....	188
APP MOVEF[ILE].....	188
APP DELETE.....	190
APP DELETEDF[ILE].....	190
APP PROPERTY CODEPAGE.....	191
APP RENAME.....	191
APP RENAMEF[ILE].....	192
Designating File Types for APP Commands.....	193
Output Redirection Commands	197
APP HOLD.....	199
APP HOLDDATA.....	200
APP HOLDMETA.....	200
APP FI[LEDEF].....	201
Application Metadata Commands and Catalog Metadata	201
Retrieving Basic Information.....	201
STATE.....	202
APP LIST.....	203
APP QUERY.....	204
Retrieving Extended Catalog Information.....	206
catalog/sysapps.....	206
catalog/sysfiles.....	207
APP HELP	209
Restricting the Use of APP Commands	209

Accessing Metadata and Procedures	211
Search Rules.....	211
Creation Rules for Procedure Files.....	212
Locating Master Files and Procedures.....	212
Accessing Existing Data Files.....	214
Creation Rules for Data Files.....	214
Data Set Names.....	219
Allocating Temporary Files	220
Temporary Space Usage and Location	223
Temporary Disk Space Usage for Non-PDS Deployment	223
Application Tools	226
EX Procedure and Amper Variables.....	226
EX EDAMAIL.....	227
4. Managing Adapters	233
Changing Common Adapter Settings	233
Configuring an Adapter	234
Changing the Adapter Configuration	236
Configuring a Remote TIBCO WebFOCUS Reporting Server	238
Complete Adapter List	238
SQL Automatic Passthru Status	239
Determining SQL Optimization	240
Data Type Support Report	242
Displaying License Counts	243
5. Metadata	245
Creating Synonyms	245
Testing Synonyms	248
Managing Synonyms	249
LNGPREP Utility: Preparing Metadata Language Files	252
LNGPREP Modes.....	254
Uploading Files	256
Preparing Spreadsheets for Upload.....	257
Starting the Upload Process.....	261

Classifying Metadata During Upload.....	262
Prerequisites for Metadata Classification.....	263
Getting Data	267
Using Data Assist	270
Introduction to Metadata.....	270
Accessing Data.....	271
Working With Data Adapters.....	272
Generating Synonyms.....	273
Opening the Synonym Editor.....	274
Understanding the Synonym Editor User Interface.....	275
Understanding Synonym Editor Frames.....	279
Enhancing a Synonym.....	296
Assigning Dimension View (DV) Roles.....	297
Creating a Virtual Field (DEFINE) or Calculated Value (COMPUTE).....	299
Using the Expression Calculator.....	301
Adding a Global Variable to a Synonym.....	306
Creating a Join.....	308
Editing a Join by Right-Clicking a Synonym Icon.....	312
Editing a Join by Right-Clicking a Connector Line.....	312
Editing a Parent Link.....	313
Adding Data Source Security (DBA).....	314
Pivoting Columns.....	319
Other Synonym Enhancements.....	321
Key Analysis Using Sample Data.....	338
6. Working With Data Flows	345
Creating a Data Flow	345
Showing the Confidence Level for Adding a Union to a Data Flow.....	351
Using the Union Editor.....	352
Preparing Data With a Flow	376
Generating Sample Files.....	377
Enabling Sampling for a Data Flow.....	378
Using Stratified Sampling.....	380

Prioritizing Sampling on Joined Tables.....	382
Editing the Fields in a Flow.....	383
Showing Profiling for a Flow.....	384
Adding or Replacing Fields.....	386
Join Editing.....	402
Applying a Filter.....	406
Loading the Data to a Target Data Store.....	408
Editing Data Flow Options.....	412
Assigning Slowly Changing Dimension (SCD) Columns.....	420
Generating Data Flows Using a Template Flow.....	426
Controlling Format Conversion Error Processing in a Data Flow.....	443
Migrating Data Flows.....	445
Viewing Checkpoint Files Generated by a Change Data Capture (CDC) Process.....	449
7. Stored Procedures.....	455
Working With Stored Procedures.....	455
Creating a Stored Procedure.....	456
Using the Reporting Server Browser Interface Search.....	457
Sending Stored Procedure Report Output to Server-Side Printers.....	459
Editing a Stored Procedure.....	459
Text Editor Menu Bar Functions.....	460
Text Editor Color Coding.....	461
Searching for Text or a Regular Expression.....	461
Configuring Text Editor Options.....	463
Using Code Folding.....	464
Using Custom Folding.....	466
Using Selection Folding.....	467
Text Editor Autocompletion.....	467
Using Tab Key Navigation.....	468
Comparing or Merging Files.....	469
Running a Stored Procedure.....	471
Running a Procedure With Different Variables.....	472
Executing a Repository FOCEXEC From EDASPROF.....	473

Prompting for Amper Variables in Procedures.....	473
Debugging a Stored Procedure	473
Using the Dialogue Manager Debugger.....	474
Setting Breakpoints.....	475
Watching Variables.....	476
Debugging Steps.....	477
Editing a FOCEXEC or Master File in the Debugger.....	478
Dialogue Manager Debugger Sample Session.....	480
Running a Procedure With Advanced Options	487
Scheduling a Procedure	488
Sending an E-mail Notification for a Procedure	493
Procedure Run-Time Processing Essentials	495
Commands in a Stored Procedure.....	495
Dialogue Manager Commands and Processing.....	496
Dialogue Manager Command Processing in a Procedure.....	498
Using Variables.....	501
Naming Conventions.....	501
Local Variables.....	502
Global Variables.....	503
Predefined System Variables.....	504
Supplying Values for Variables.....	509
General Rules.....	509
Supplying Values in the EXEC Command.....	509
Displaying Execution Flow.....	511
-DEFAULT[S] Command.....	511
-SET Command.....	512
Branching.....	513
Screening Values With -IF Tests.....	515
Calling Another Procedure.....	516
Nesting.....	517
The EXEC Command.....	517
ON TABLE HOLD.....	518
ON TABLE PCHOLD.....	519

Creating Expressions.	519
Arithmetic Expressions.	520
Alphanumeric Expressions.	521
Logical Expressions.	522
Compound Expressions.	525
ON TABLE HOLD.	526
ON TABLE PCHOLD.	527
Platform-Specific Commands and Features	528
DYNAM Command (z/OS PDS and HFS Deployments).	528
Use of Data Sets.	530
The ALLOCATE Subcommand.	531
The CONCAT Subcommand.	540
The FREE Subcommand.	541
The CLOSE Subcommand.	542
The COPY Subcommand.	543
The COPYDD Subcommand.	545
The DELETE Subcommand.	546
The RENAME Subcommand.	547
The SUBMIT Subcommand.	548
The COMPRESS Subcommand.	549
Comparison of TSO Commands, JCL, and DYNAM.	549
DDNAME Translation in User-Written Subroutines: MSODDX.	551
MSODDX Calling Environment and Migration From an MVS Server.	551
Using MSODDX in a Typical Application.	552
MSODDX Parameters.	552
ASMDDX: Sample Assembler Exit for MSODDX.	554
COBDDX: Sample COBOL Exit for MSODDX.	556
FILEDEF Command Under UNIX, Windows, IBM i, and z/OS.	558
Other FILEDEF Features.	560
OFFLINE Printing.	561
UNIX and z/OS HFS Deployment.	562
z/OS PDS Deployment.	562
Windows.	563

IBM i.....	563
OFFLINE to DISK Versus PRINTER.....	564
Other Printing Information.....	564
Stress Testing a Procedure	565
Recording User Actions Into a Script Using Run Stress.....	566
Running Impact Analysis Reports	569
8. Workspace Manager	573
Configuring Workspace Manager	573
Workspace Manager Configuration Privileges.....	574
Configuring an Agent Service.....	576
Configuring Deployment Modes.....	580
Notification	582
EVENT_TEXT and EVENT_ID.....	586
Login Message.....	587
Workspace Configuration Settings	589
Reporting Server Web Console Appearance.....	590
Locale Settings.....	591
Displaying Settings for Reporting Server Browser Interface Preferences.....	601
TIBCO FOCUS Sets and Info.....	605
Core Engine Settings.....	605
Core Engine System Variables.....	608
Core Engine Global Variables.....	609
Displaying HOLD Keywords.....	610
Displaying FOCUS Limits.....	611
Displaying Error Messages.....	613
Displaying the Currency Report.....	614
Determining Function Parameters.....	615
Workspace SET.....	617
Configuring Geographic Information.....	627
Editing the Geographic Configuration.....	634
Migrate Settings.....	658
Monitoring TIBCO WebFOCUS Reporting Server Activity	659

Monitoring TIBCO WebFOCUS Reporting Server Status.	659
Monitoring Data Services.	660
Displaying and Customizing a Data Service Agents Page.	661
Displaying Status Statistics and Aggregated Statistics for Data Services.	669
Working With the Connection Queue.	674
Performing Administrative Tasks for Data Services.	676
Monitoring Java Services.	681
Displaying and Customizing a Performance Report for Java Service Agents.	681
Displaying Aggregated Statistics for Java Services.	686
Performing Administrative Tasks for Java Services.	687
Monitoring Special Services and Listeners.	689
Setting a Web Session Timeout Warning.	695
Deferred Lists.	695
Web Sessions.	697
Performing Administrative Tasks for Special Services and Listeners.	700
Editing Configuration Files	701
Configuring Java Services	702
Fine-Tuning the TIBCO WebFOCUS Reporting Server.	704
Cluster Manager	707
Adding Auto Cluster Manager Self Registration to a Remote TIBCO WebFOCUS Reporting Server.	710
Cluster Manager Global Monitoring Console.	721
Resetting Cluster Manager Statistics.	725
9. Listeners and Special Services	727
Using the Listeners and Special Services Configuration Panes	727
Special Services	727
Running and Configuring the FOCUS Database and FOCUS Database Server.	728
The Scheduler and Deferred Service.	731
Forcing a Scheduler Scan.	744
Storing Deferred Reports in an Application Directory or an SQL Repository.	744
Listeners	745
Configuring Listeners.	745

Other Listener Options.	754
10. Troubleshooting	757
Viewing Version Information	757
Analyzing TIBCO WebFOCUS Reporting Server Activity	759
Analyzing FOCUS Database Server Activity	766
Controlling the HLIPRINT Transaction Log.	766
Controlling HLIPRINT Output.	767
Tracing TIBCO WebFOCUS Reporting Server Activity	768
Configuring Traces.	769
Filtering Traces.	772
Showing the Session Log.	773
Changing the Size of the Session Log Buffer.	776
z/OS-Specific Troubleshooting for the Unified TIBCO WebFOCUS Reporting Server	777
Traces and Saved Diagnostics.	777
U4039 Abends.	778
System Dumps.	778
Freeing Data Sets Allocated to the Unified TIBCO WebFOCUS Reporting Server.	779
z/OS-Specific Troubleshooting for z/OS HFS Deployment	780
Recovering Orphaned Shared Memory Segments on z/OS.	780
INSUFFICIENT AUTHORITY TO GETPSENT Messages on z/OS.	781
Recording and Reproducing User Actions	781
Troubleshooting the Reporting Server Browser Interface	792
TIBCO WebFOCUS Reporting Server Processes	792
Server Work Space Manager Daemon.	792
Deferred Requests Special Services Daemon.	792
Server Log Daemon.	793
Server Check Up Daemon.	793
HTTP Listener Daemon.	793
TCP Listener Daemon.	793
FDS Special Services Daemon.	793
Java Services Daemon.	793
Agent Daemons.	793

Gathering Diagnostic Information for Customer Support Services	794
A. Starting and Stopping the TIBCO WebFOCUS Reporting Server	803
Using TIBCO WebFOCUS Reporting Server Management Command Options	803
TIBCO WebFOCUS Reporting Server Environment Variables	813
B. Unicode Support	819
Unicode and the TIBCO WebFOCUS Reporting Server	819
Accessing Unicode Data	821
Selecting, Reformatting, and Manipulating Characters	828
Sort Order Under Unicode	830
Added Unicode Support for Master Files, Data Files, and Application Directory Names	830
Unicode PDF Output	831
C. EDATEMP and NFS-Mounted Disks	833
EDATEMP and NFS-Mounted Disks	833
D. Caching Support for Node.js	835
Node.js Installation and Configuration	835
Node.js Process	836
Node.js Prerequisites	836
E. Platform-Specific Methods for Specifying EDAEXTSEC	843
Platform-Specific Methods for Setting EDAEXTSEC	843
Legal and Third-Party Notices	845

The Reporting Server browser interface enables you to configure, operate, monitor, tune, and troubleshoot the WebFOCUS® Reporting Server environment from a single, easy-to-use interface.

In this chapter:

- [Using the Reporting Server Browser Interface](#)
 - [Menu Options](#)
 - [Configuration Overview](#)
 - [TIBCO WebFOCUS Reporting Server Profiles](#)
 - [Profile Commands](#)
-

Using the Reporting Server Browser Interface

The Reporting Server browser interface enables you to remotely view and manage the server environment. From a single, easy-to-use interface, you can:

- Select, add, and configure data adapters.
- Create and manage adapter metadata.
- Control data access and security.
- Configure remote servers.
- Configure, edit, and run applications and deferred query processing.
- Configure communications and special services.
- Operate, monitor, tune, and troubleshoot your server.
- Edit configuration files.
- Migrate from a previous server release.
- Control aspects of your own environment.

Note: Server management commands are also available for starting, stopping, and managing the server. For more information, see [Starting and Stopping the TIBCO WebFOCUS Reporting Server](#) on page 803.

Procedure: How to Open the Reporting Server Browser Interface

To open the Reporting Server browser interface:

1. If the server is not running, start it using the instructions for your platform in the installation guide.
2. Type the following URL in the address bar of the web browser:

`http://ip_address:http_service`

where:

`ip_address`

Is the IP address of the machine on which the server is running.

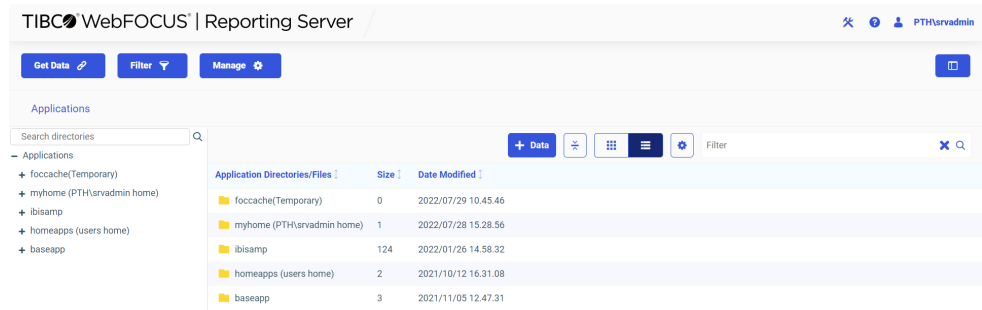
`http_service`

Is the value for HTTP Service (port number) entered during the server configuration procedure.

If you have trouble contacting the server, see [Troubleshooting the Reporting Server Browser Interface](#) on page 792.

If you are running your server with security, see [Security Providers](#) on page 36 for information on the requirements and options associated with the security mode you are using.


The Reporting Server browser interface home page opens, as shown in the following image.



When the Reporting Server browser interface opens to the Applications page, only the top-level folders are displayed, and when an application folder is selected, only files in that folder are shown.

If you select an application folder that has no files, a message telling you to select a different application displays on the empty canvas, as shown in the following image.



The Tools menu () provides access to the different features of the Reporting Server browser interface.

The browser cache files for each server build will be stored separately. This means that multiple Reporting Server browser interface sessions using the same browser will not impact each other. Therefore, you will not have to clear the browser cache to switch between multiple servers using the same machine and browser.

To implement separate cache files, each Reporting Server browser interface URL will automatically be invoked using a unique identifier. For example, if you enter the following URL:

<http://mypcname:8021/wc/wcstart.html>

The URL will automatically be issued with a unique identifier. For example:

<http://mypcname:8021/A4643588A8A7C435/wc/wcstart.html>

All other URLs generated by web connections to the server will use the same identification string.

3. Select an option from the Tools menu to access the corresponding page and ribbon. For tasks that display information or require no additional navigation, the page displays corresponding information in the right pane. For tasks that require additional task-centered navigation, the page opens a new panel or window. Use the resources pane of the new window to access the corresponding information.

On Reporting Server browser interface resources panes:

- Right-click behavior opens pop-up menus, by default.
- Double-click behavior executes the first option on a pop-up menu.

The Reporting Server browser interface offers online help with two kinds of links to appropriate sections of this documentation:

- Click the ? icon on the menu bar and choose *Online Server Help* to display the full Reporting Server browser interface help system.
- Click the ? icon, when available, to display contextual help specific to the item.

Reporting Server Browser Interface Ribbon

Reporting Server browser interface functionality is available through a ribbon at the top of the page.

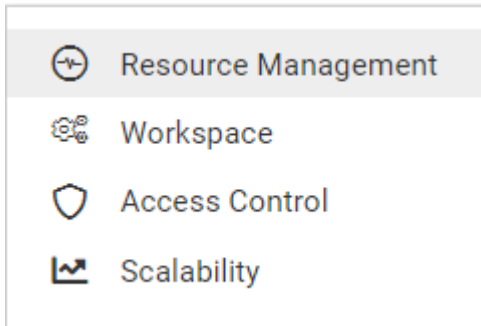
The ribbon available from the Applications page is shown in the following image.



The interface includes the following menus at the top right of the page, as shown in the following image.



- ❑ The Tools menu provides access to other Reporting Server browser interface pages, as shown in the following image.



To return to the Applications page, click the WebFOCUS Server icon at the top left of the page:



- ❑ The help menu opens the help content.
- ❑ A menu under the user ID of the connected user is displayed, if the server is running with security on. If it is not running with security on, a User menu appears.

Reporting Server Browser Interface Hover Menu

Pages and panes in the Reporting Server browser interface can be rearranged, resized, or closed using the Hover menu. The Hover menu appears when you point your cursor to the top right corner of a pane, shown in the following image.



This menu contains the following options, some of which depend on the page that is open:

- ❑ **Drag and drop.** Allows you to move the page within the Reporting Server browser interface. A shaded area indicates where the page will be placed.

You can also access this option by pointing to the top of the right pane. The cursor becomes a four-headed arrow, indicating that you can move the page.

- Menu.** The menu contains options to control the layout of the pane of the Reporting Server browser interface. The same menu appears when you click the active tab or right-click a tab. The following options are available:
 - Maximize.** The selected tab fills the Reporting Server browser interface.
 - Close.** Closes the selected tab.
 - Layout Options.** Opens the Layout Options dialog box. The Layout Options dialog box provides the following options:
 - Hover-tab.** Allows you to select whether the hover menu should be above the content or on top of the content.
 - Frame bar.** Allows you to select whether to show or hide a small pane that contains icons that let you show or hide each frame on the page.

You can always show the frame bar, always hide the frame bar, or show only hidden frames and nested layouts, which is the default.
 - Max content per frame.** Allows you to specify the maximum number of tabs to show when the View maximum set option is selected.
 - Show Layout Tree.** Shows the properties of the container and frames.
 - Maximize.** The selected tab fills the Reporting Server browser interface. Once it is maximized, a *Restore* button is added to return to the original view.
 - Hide.** Closes the selected tab.

Menu Options

Each Console page provides options on the ribbon and right-click options from the resources pane.

The following options are available on the right portion of the page.

User Menu Options

The following options are available from the User menu on the right portion of the page.

End Current Session or Sign Out

This option signs the user out and opens the Reporting Server browser interface Sign-in page.

My Console

These options provide users with functionality that supports the tasks they are authorized to perform. The options that a user has under My Console depend on the user role and general privileges assigned to the user.

In the descriptions that follow, each option applies to the user who is connected to the current Reporting Server browser interface session.

Show My General Privileges. Shows the privileges set by the Server Administrator for the user who is logged in to the current Reporting Server browser interface session. Privileges can be customized for each role, group, or user. See [General Privileges](#) on page 77, for details on the information supplied on the General Privileges pane.

Change My Password. Enables you to specify a new password to replace your current password. When you click this console option, the dialog box shown in the following image is displayed. Supply the requested information, and click *Update*.

Note: The LDAP security mode also allows you to change your password in the LDAP directory.

The image shows a dialog box titled "User Information". It contains the following fields and options:

- User: **pgmtst1**
- Current Password: [password field with 5 dots]
- New Password: [password field with 5 dots]
- Confirm Password: [password field with 5 dots]
- Update LDAP password
- Update password in admin.cfg file (needed for scheduled runs ONLY)
- Buttons: Update, Cancel

Manage My Agents. Enables you to monitor and manage Data Service Agents, based on your role and privileges. For example, if you are a Basic user, by default you can monitor and manage your own agents (that is, the agents that match your user ID). If granted the applicable privileges, you can also monitor, or monitor and manage, the agents of the other users in your group (that is, users with the same group ID).

When you click *Manage My Agents*, a Performance Report for Data Service Agents is displayed, showing statistics for the agents that you are authorized to see. You can customize and filter the report to display only those statistics that you are interested in.

To determine your privileges, click *Show General Privileges*. See [General Privileges](#) on page 77 for details on the information supplied on the General Privileges pane.

Save My Last Report. Saves the last report (procedure) that you ran with the values that you chose for the amper variables. The report is saved as a new, separate procedure. You can run the saved report without having to manually supply the values for the amper variables as you did before.

When you click this console option, you can select the application directory in which to save the report and enter a file name for the report.

Edit My Profile. Enables you to make changes to your current user profile settings.

Login Info

This option displays user login status and browser-related information. Login status information includes user name, group, and role, location of all profiles (user, group, and role) if they exist, and type of authentication (explicit or cookie). HTTP information reflects the properties of the current browser session, for example, the cookie.

Enable Traces

This option enables traces to provide diagnostics assistance. For more information, see [Tracing TIBCO WebFOCUS Reporting Server Activity](#) on page 768.

Clear Session Log

This option deletes all lines from the session log.

Session Log

This option opens the Session Log window, which provides a record of server configuration and activity.

Export for Printing

If a report, such as a Data Services Agents report, is open in the Reporting Server browser interface, clicking this option opens the report in a web browser.

Preferences

This option enables you to indicate preferences for Run Options, Web Console Appearance, Reporting Options, and Diagnostic Options, as shown in the following image.

Run Options are:

- Default Data Service** The choices for Default Data Service are DEFAULT, WF_DEFAULT, SCHED_DEFAULT, and DFM_DEFAULT.
- Maximum number of seconds before the popup timer appears**

The Web Console Appearance options are:

- Ribbon's Text.** The choices are Adhere to server appearance default, Icon Hide, and Show. The default is Hide.
- Report Style.** The choices are Warm, Standard, and Flat. The default is Warm.
- Default page on User Login.** The choices are Applications, Adapters, Resource Management, Workspace/Monitor, Access Control, and Scalability. The default is Applications.

- Show warning when leaving page.** The choices are Yes and No. The default is Yes.

The Reporting Options are:

- Maximum number of rows for test reports.** The default is 50.
- Maximum number of columns for test reports.** The default is 999,999.
- Maximum number of identical error messages.** The default is 10.
- Omit missing (null) values in Data Profiling charts.** The choices are Yes and No. The default is No.

The Diagnostic Options are:

- Enable recording.** The choices are Yes and No. The default is No.

Indicate your preferences, and click *Update*.

Language

This option controls the language in which the Reporting Server browser interface is displayed. For example, all menu options are translated into the selected language. This console option is available if the server is configured for National Language Support (NLS).

Configuration Overview

The characteristics of an individual server are determined by a set of configuration files. These configuration files define the protocols, services, and data sources supported by the server. In the UNIX, z/OS HFS Deployment (USS), Windows, and IBM i (formerly known as i5/OS) environments, these files are maintained in directories designated by the environment variable EDACONF.

The initial installation procedure creates an installation instance, represented as the logical name EDAHOME, and one default configuration instance, known as EDACONF. After you successfully install the server, you can configure a default configuration instance to create an operational instance, or create an additional one at your site. An operational instance of the server is one that is configured to support specific protocols, services, and access to data sources.

You must run the Installation/Configuration Utility for each new instance of the server that you want to configure, and then use the Reporting Server browser interface to configure all necessary functionality.

To configure a server instance, use the Reporting Server browser interface to:

1. Select and configure data adapters.

2. Optionally, configure remote servers.
3. Configure communications nodes and protocols.
4. Optionally, set parameters for deferred query processing.

If necessary, you can edit configuration files. Some configuration errors can make the server start in a limited mode called safe mode, in which the Reporting Server browser interface is still operational, to enable you to correct the errors.

TIBCO WebFOCUS Reporting Server Profiles

A server or global profile, `edasprof.prf`, is created during installation. You can customize this profile, which is applied to all users. Optionally, you can create separate profiles to customize the server environment for a role, a group of users, or an individual user. You can also create service profiles, which specify settings for the server environment. For information, see [Profile Settings](#) on page 618. Profiles are executed in the following sequence: server, service, role, group, and user.

A profile can include almost any command that a client application can send to the server, however, profiles are used most frequently for application set up commands, such as SETs, search path set up, and DBMS connection information, which may vary by adapter.

- ❑ For information about the syntax for commands typically used in profiles, see [Profile Settings](#) on page 618.
- ❑ For a description of the profile commands that are available for a particular adapter, go to the *TIBCO WebFOCUS® Adapter Administration* manual. In the adapter chapter, look for topics on customization and optimization.
- ❑ For information about the relationship between profiles and security modes, see [Security Providers](#) on page 36.

Two key aspects of profile customization are level and associated search order. The server processes all profiles found in the search order. If it finds duplicate settings or commands, the last setting or command processed will be active for the connection.

When you create a profile, take into account that the more processing performed by the profile, the more time it takes for an application to connect to the server. An exception is pooled deployment, where only one profile is processed when the server agent initially starts or is refreshed.

Profile Level

The server supports various levels of profiles to provide flexibility for designing and running production applications. The order in which the following profiles are listed correlates with the search order in which they are processed by the server when multiple profiles exist.

- ❑ **Global profile.** The first level of profile, the global profile, is a startup file that is automatically created during installation and configuration of the server. It contains default environment settings required for the correct operation of the server.

The global profile remains in effect throughout a user session. You can modify the global profile default settings. You can also add any commands or code that all connected users require before application processing begins.


Characteristic	Description
Name	edasprof.prf
Location	UNIX: \$EDACONF/etc directory Windows: \$EDACONF\etc directory IBM i: IFS \$EDACONF/etc directory
Override	Set the environment variable EDAPROF to point to an alternate location of edasprof.prf (you cannot rename edasprof.prf).
Sequence in search order	First
Scope	Applies to all connected users.

- ❑ **Service profile.** A service profile specifies settings for the server environment, but the settings in this level of profile apply only to users associated with a specific service. When a user connects to the server, the service profile settings are applied and remain in effect throughout the user session.

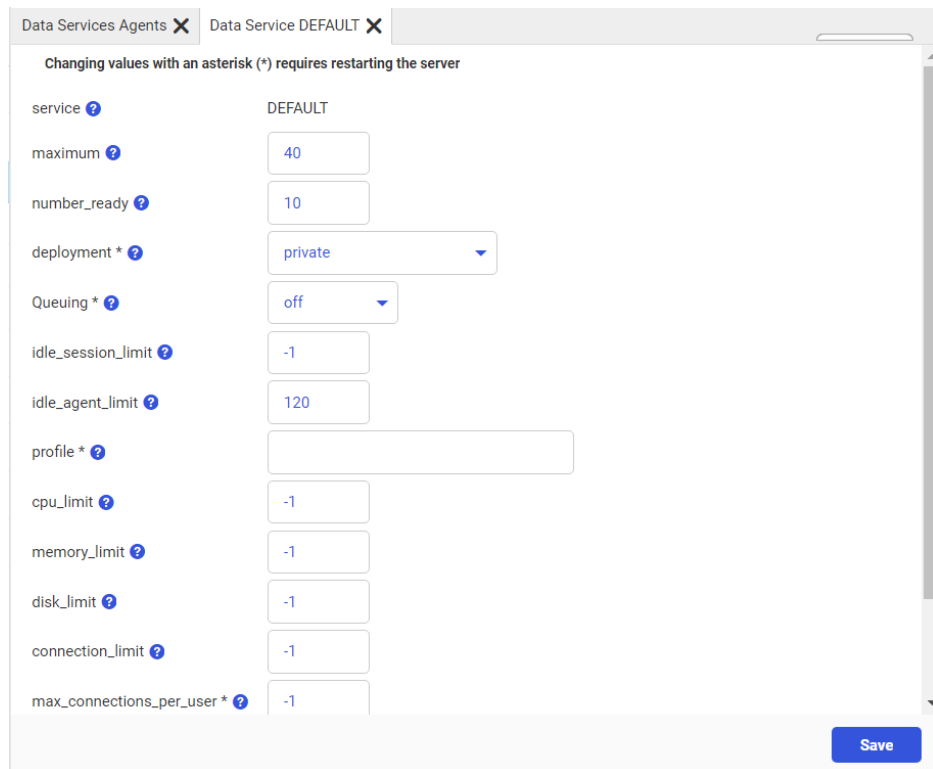
A service profile is a procedure (.fex) that will be executed on an agent startup for that service.

Create a service profile by:

1. Creating a new procedure in a folder that is on the APP PATH of the server.
2. Enter the settings you want to apply to the service.
3. Save the procedure with the name of the service, for example, default.fex.
4. On the Workspace page of the Reporting Server browser interface, right-click the service and click Properties, as shown in the following image.



The Properties page for that service opens, as shown in the following image.



5. Enter the *app/profile_name.fex* in the profile entry field, and click Save.

A dialog box opens saying the server needs to be restarted.

6. Click OK to restart the server.

When the Workspace restarts, the service profile procedure is run.

Characteristic	Description
Name	<i>app/name.fex</i>
Location	Follows the normal search path for procedures (FOCEXECs).
Override	None
Sequence in search order	Second

Characteristic	Description
Scope	Applies to all users of the same service.

- ❑ **Role profile.** A role profile specifies settings for the server environment, but the settings in this level of profile apply only to users assigned to a specific role. When a user connects to the server, the role profile settings are applied and remain in effect throughout the user session.
- ❑ **Group profile.** A group profile specifies settings for the server environment, but the settings in this level of profile apply only to users associated with a specific security group. When a user connects to the server, the group profile settings are applied and remain in effect throughout the user session.

A group profile may contain settings that are the same as those in a global profile.

Server Administrators need to register group profiles and users to give them privileges that are different from those they inherit from their default role. Group profiles are used to set DBMS connections, variables, and other settings that should be executed for group users when they connect.

If you use an editor, create the group profile in the correct location.

Characteristic	Description
Name	<i>group_name.prf</i>
Location	UNIX: <i>base_directory_name/ibi/profiles</i> Windows: <i>base_directory_name\ibi\profiles</i> IBM i: IFS <i>base_directory_name/ibi/profiles</i> directory
Override	Set the <i>edaserve.cfg</i> keyword <i>edaprfu</i> to point to an alternate location of <i>group_name.prf</i> .
Sequence in search order	Third
Scope	Applies to all users of the same security group.

- ❑ **User profile.** A user profile specifies settings for the server environment, but the settings in this level of profile apply only to a specific user ID. When a user connects to the server, the user profile settings are applied and remain in effect throughout the user session.

A user profile may contain settings that are the same as those in a global profile.

Procedure: How to Create a User, Group, or Role Profile

You can create a user profile at any time by:

1. Go to the *Workspace* page.
2. Expand the *Configuration Files* folder in the resources pane.
3. Right-click the *User/Group Profiles* folder, and click *New Profile*.

The Edit New Profile page opens.

4. Enter the profile parameters and click the *Save As* button.
5. Enter a profile name and click *Save*.

You can also use any standard system editor. If you use an editor, create the profile in the correct location.

Characteristic	Description
Name	<i>user_id.prf</i> , <i>group_id.prf</i> , or <i>role_id.prf</i> .
Location	UNIX: <i>base_directory_name/ibi/profiles</i> Windows: <i>base_directory_name\ibi\profiles</i> IBM i: IFS <i>base_directory_name/ibi/profiles</i> directory
Override	Set the <i>edaserve.cfg</i> keyword <i>edaprfu</i> to point to an alternate location of <i>user_id.prf</i> .
Sequence in search order	Depends on which type of profile and how many other profiles apply to the same user, group, or role.
Scope	Applies to a specific user, group, or role.

Profile Commands

The following topics describe commands that you can include in any of the supported server profiles. These commands affect the behavior of the server for the duration of the connected session. You can code additional commands, such as data access commands, in any of the supported server profiles. For more information, see specific adapter topics in the online help.

For more information on the server profile commands, launch the online help using the adjacent question mark (?) icons in the Reporting Server browser interface.

Profile Command Formats

Each server command takes one of three possible formats. The syntax and an example of these formats follow. Ensure that you use the correct format for any server command you use.

Note: The list of commands is available from the *Settings* button on the Reporting Server browser interface Workspace page. We recommend that you use the Reporting Server browser interface settings panes to add the desired settings to your chosen profiles.

Syntax: How to Use a Direct SET Command

To use a direct SET command

```
SET command=value
```

where:

```
command
```

Is the server command.

```
value
```

Is the value selected from the available choices.

Example: Using a Direct SET Command

```
SET SQLENGINE=DB2
```

Syntax: How to Use an SQL Engine SET Command

To use the SQL Engine SET command

```
ENGINE sqlengine SET command value
```

where:

```
sqlengine
```

Is the target DBMS. You can omit this parameter value if you previously issued the SET SQLENGINE command.

```
command
```

Is the server command.

value

Is the value selected from the available choices.

Example: Using an SQL Engine SET Command

```
ENGINE SQLORA SET OWNERID EDAUSER
```

Syntax: How to Use an SQL Translator Command

To use an SQL Translator command

```
SQL  
SET command=value  
END
```

where:

command

Is the SQL Translator command.

value

Is the value selected from the available choices.

Example: Using an SQL Translator Command

```
SQL  
SET APT=OFF  
END
```



Chapter 2

TIBCO WebFOCUS Reporting Server Security

This chapter describes the security providers that can be configured on a server. It also describes how an administrator enables a particular security provider and assigns privileges to users and groups.

In this chapter:

- [Security Overview](#)
 - [Configuring Authentication](#)
 - [Configuring Privileges and Other Authorizations](#)
 - [Configuring Encryption](#)
 - [Access to Connection Information in TIBCO WebFOCUS Procedures](#)
 - [Configuring Multi-Tenant Deployment](#)
-

Security Overview

Server security is based on a three-step process: authenticating the connecting user, establishing the group and role for the user, and then assigning privileges and other authorizations based on the user, group, and role.

Note: NLS characters are supported for user ID, group name, password, and domain name on the Reporting Server browser interface Access Control configuration pages.

The Access Control page on the Reporting Server browser interface provides a server administrator with the tools needed to manage server security. Using ribbon buttons and right-click options, the server administrator can perform the following security operations:

- Manage security providers.
- Edit Access Control and Encryption settings.
- Filter the Access Control tree display.
- Register users, groups, and roles.
- Show privileges for specific users, groups, and roles.

Authentication

Authentication can be performed based on the following types of credentials:

- ❑ **Explicit.** User ID and password are supplied when signing in to the server.
- ❑ **SSO (Single sign on).** The user is authenticated by a single sign on product prior to signing in to the server, and a token is passed to the server. Single sign on products include Kerberos, NTLM tickets, or browser cookies, such as passed by MYSAPSSO2.
- ❑ **Trusted.** The user ID is submitted without the password from the application server or mid-tier, which has the user already authenticated.

Security Providers

A security provider is used to authenticate the incoming connection. When a security provider is configured, it is added to the `edaserve.cfg` file as a security provider block. Multiple active security providers can be configured, in which case the `security_provider` attribute will consist of a comma-separated (,) list of security providers. The first one on the list is the primary provider.

The security provider can be of several types:

- ❑ **OPSYS.** Users are authenticated against the native OS security repository, such as the windows domain controller, z/OS RACF/ACF2, UNIX password files, or other operating system security products using the native OS API.
- ❑ **LDAP.** Users are authenticated against an LDAP server or active directory using the LDAP API.
- ❑ **DBMS.** Users are authenticated by connecting to a DBMS such as Oracle, SAP, or Db2.
- ❑ **PTH.** Users are authenticated against a list of users kept internally by the server in the `admin.cfg` configuration file.
- ❑ **CUSTOM.** Users are authenticated using a custom procedure.

After the server is installed, the default security provider configured by the installation process is the PTH (server internal) provider. This provider keeps a list of groups and group membership in the `admin.cfg` file. During installation of the server, a server administrator user ID is created in the `admin.cfg` file. The installer can accept the default user ID (`svadmin`) or change it, but the installer must provide a password for this ID.

The server administrator can choose whether to register users and groups for the primary security provider with one-part or two-part names. With two-part names, the user ID or group name is prepended with the provider name and a backslash (\). For example:

`OPSYS\user1`

or

`LDAP01\user2`

By default, all users and groups are registered using two-part names. For the primary provider, you can configure one-part names. Secondary provider registrations always use two-part names. To configure one-part registrations for the primary security provider:

1. On the Access Control page, click *Access Control* on the *Settings* menu.
2. Select *n* from the *prepend_provider_name* drop-down list in the General section of the page, as shown in the following image.

The screenshot displays the 'General' configuration section of the TIBCO WebFOCUS Reporting Server Administration console. The 'prepend_provider_name' dropdown menu is set to 'y'. Below this, the 'Web Console Anonymous Access' section contains three empty input fields for 'anonymous_id', 'anonymous_pass', and 'confirm password'. The 'Password' section at the bottom has a 'password_change_delimiter' set to '.' and a 'password_change_wlogin' dropdown set to 'n'. An 'Apply and Restart Server' button is located at the bottom right of the configuration area.

3. Click *Apply and Restart Server*.

The *prepend_provider_name=n* attribute is added to the *edaserve.cfg* file.

If the server administrator did not configure one-part names for the primary security provider, the user ID will be created as a two-part name in the form *security_provider\user_ID*, for example, `LDAP01\user1`.

When the server starts with security PTH, the PTH\svadmin user ID (or the one configured during installation) is the server administrator user that is used to connect to the Reporting Server browser interface, and that user can make additional security changes to the server, such as adding and changing active security providers and changing access control privileges for users, groups, and roles. We recommend that you keep the PTH provider as an active security provider, so user PTH\svadmin can be used as a backup Server Administrator.

All new security subjects for all security providers (except for the primary provider, if one-part names were configured) will be registered with a two-part name, provider\userid or provider\groupid. For example, the following group named grp1 is registered under the LDAP1 security provider:

`LDAP1\grp1`

The following user, whose ID is user1, is registered under OPSYS security in the IBI domain:

`OPSYS\IBI\user1`

If one-part name registrations exist from a prior release, they are respected and are assumed to be primary provider registrations.

Note: If the silent installation is used, the server administrator user ID with its associated password needs to be provided. If none is provided, the default user ID PTH\svadmin with the password *svadmin* will be configured.

OPSYS security has specific operating requirements in order to be activated. For details, see [Configuring OPSYS Authentication](#) on page 46.

If you want to use an OPSYS, LDAP, DBMS, or CUSTOM security provider, you must configure it on the Reporting Server browser interface Access Control page. When you configure one of these security providers, the server edits the edaserve.cfg file to add a block for the new security provider. You can then use the Reporting Server browser interface to change to this provider. In response, the server sets the security_provider attribute in the edaserve.cfg file to this provider name, after which this provider is considered the primary provider and defines the security used by the server.

You can configure multiple security providers, and the server will add multiple provider blocks in the edaserve.cfg file. When multiple providers are chosen to be active, the security_provider attribute in the edaserve.cfg file is set to the list of active providers. In this case, the first provider on the list is considered to be the primary provider and defines the security used by the server.

Note: In order for a signed-in user to start or manage a server using the edastart command script, the signed-in user ID has to be registered as a server administrator in the admin.cfg file, even if OPSYS is not an active security provider. This is required in order to protect the server from unauthorized actions by users who can sign in to the system where the server resides. If one of the active providers is OPSYS, the server must be authorized to run security OPSYS and it will start with security OPSYS.

EDAEXTSEC can be set to OFF in order to override the security provider set in the security_provider keyword of the edaserve.cfg file and start the server with security OFF. For information on setting EDAEXTSEC in different operating environments, see [Platform-Specific Methods for Specifying EDAEXTSEC](#) on page 843.

An example of configuring multiple providers is having two distinct LDAP providers configured against two separate LDAP servers that can contain different users and different groups.

One provider is considered the primary provider. For all providers, the user name is a two-part name prefixed with the provider name, for example, LDAP1\usera. Note that users (or groups) from two providers that have identical names (such as LDAP1\usera and LDAP2\usera) are considered different security subjects and can be given different privileges and profiles.

At sign-in time, a user needs to select the security provider name in addition to the user ID and password. If a user signs in to the server using the primary provider, the server will automatically prepend the user ID with the provider name, if the user does not supply a two-part name.

OPSYS providers can be configured with other security providers. On Windows, multiple Windows domains act as different security providers and users. For example, domain1\usera and domain2\usera are not identical.

Privileges and Other Authorizations

The server supports a large number of objects, and access to them should normally be restricted.

DBMS

With DBMS authorization, the main object to be protected is the DBMS-resident data that the server reads and writes. Access control is implemented using ENGINE CONNECTION_ATTRIBUTES statements that define the security attributes that are used by the server agent for connecting to the DBMS. The connection security type depends on the DBMS and can be one of three subtypes:

- Explicit.** The userid/password is stored in a profile (encrypted).

- ❑ **Password pass-through.** The userid/password combination or cookie, such as MYSAPSS02, that were provided during the connection to the server is passed to the DBMS connection.
- ❑ **Trusted.** This type of connection applies to security provider OPSYS where the operating system process of the agent impersonates the user, and the DBMS connection is derived from that process.

The CONNECTION_ATTRIBUTES statement itself can be defined in the user, group, role or server profile, and they override each other in this order.

It is common initially to create connections in the edasprof server profile. They will then be inherited by all users. Subsequent configuration steps may define connections on other levels, such as the group or user level.

The CONNECTION_ATTRIBUTES statement in effect passes the credentials to the DBMS, and the DBMS server ensures the correct access control: read and read/write rules on the DBMS tables, views, columns, and rows.

Server DBA

The server engine can add an additional level of access control by defining DBA rules in synonyms to restrict access to columns and certain data values. This applies only to data access using the WebFOCUS language, and not to Direct SQL Passthru requests. For more information about Server DBA, see the *Describing Data With TIBCO WebFOCUS® Language* manual.

Data and Application Folder Access

The server reads operating system files using standard I/O calls.

With the OPSYS security provider, the proper OS read/write privileges should be given using native OS tools such as RACF rules or Windows/UNIX permissions. Run-time access control is achieved by the operating system validating the data agent access to any given file based on the agent process impersonation of the connecting user.

To protect files in other modes (LDAP, PTH, DBMS), the server access control feature should be used to assign read/write privileges to OS files and folders. The privileges are managed from the Reporting Server browser interface Access Control page and stored in the server admin.cfg file. They can be assigned at the role, group, and user levels and be allotted on the folder and file levels. Normal inheritance rules apply. A subfolder inherits from its parent folder, a file inherits from its parent folder, a user inherits from the group, and so on.

Note that this protection applies to WebFOCUS metadata objects in application folders such as synonyms, FOCEXECs, HTML files, style files, and to data files such as FOCUS data sources and sequential files.

Application folders and files can have four types of privileges, read, write, execute, and list. For data files such as FTM files, only two privileges are used, read and write.

You can disable operating system shell commands using the SET OPSYSCMD command or the NOSYS general privilege. You may want to consider disabling these commands as they are not subject to access control.

General Server and Reporting Server Browser Interface Actions

The Reporting Server browser interface provides various actions, such as adapter configuration, server configuration and monitoring, Resource Analyzer configuration, metadata creation, and other actions, such as the ability to issue Direct SQL Passthru requests. To manage access to these actions, the server administrator uses the server access control feature to assign *general* privileges to users, groups, and roles.

Groups

Groups are collections of users typically defined by the security provider for the purpose of giving common authorization to multiple, instead of individual users. Groups are defined either on an external repository such as RACF or LDAP, or internal PTH. The DBMS security provider does not support groups. Except for PTH, user provisioning and all associated administration tasks (such as password reset, invalidating users, and volume user creation) are accomplished using the respective security provider software and are outside of the Reporting Server software.

Roles

Roles are a collections of privileges defined on the server for the purpose of registering groups and users to the collection (as opposed to assigning privileges to individual users). Five predefined user roles (Server Administrator, Application Administrator, Operator, Basic, and None) are created at server installation time. For example, the Server Administrator role has all privileges, and the Application Administrator role has privileges related to the creation of applications, such as creating synonyms and related tasks.

New roles can be added and existing roles can be modified to include and exclude individual privileges.

New installations are done with no access to files and directories outside of the server installation, and with no privilege to run operating system commands and direct SQL requests to a DBMS, for all roles except Server Administrators. All users have full access to all applications under APPROOT,

Registering Users and Groups Under Roles

When the server is installed, the user ID specified as the server administrator is automatically registered as Server Administrator. All other users and groups are defaulted to the server default role—*Basic*.

Note that no group, user, or role profile exists initially after the installation. It is typical at the initial stage to have the DBMS connections defined in the server profile, *edasprof*. Privileges are derived from the Basic role and are in effect for all users except the Server Administrator.

Recommended steps for securing the server:

- Review default privileges for role Basic and change them if necessary. You can change the default role to another role.
- Review and change the file privileges for the default role at the file root level.
- Customize groups:**
 - Register groups to roles. For example, register LDAP group GRP1 to the Application Administrator role.
 - Create group profiles with DBMS connections and APP PATH. For example, for group GPR1, create *grp1.prf* (through the Reporting Server browser interface Adapter page or the application path configuration).
 - Customize application level file privileges for groups. For example, group GRP1 can have read but not write access to *baseapp*.
- Optional steps:**
 - Create new roles or customize the role privilege list. This should be done only if existing roles are not satisfactory. It allows registering groups to these roles.
 - Register users. This should be done only if groups are not created adequately, and for special users such as managers.
 - Register file privileges on the file level. This should be avoided in favor of folder level privileges, as they are inherited by files and are easier to manage.

Registration of Users and Groups for Inactive Security Providers

From the Access Control page, the server administrator can register users and groups in server roles for a security provider that is not currently activated. These registrations will be available when the server is started with security ON and with that provider activated.

In addition, a server running with security OFF has an Access Control page that a server administrator can use to register users and groups and to activate and deactivate providers.

One use for this feature is to register a Server Administrator for a Provider that will be active after the next server restart.

Encryption

Using the Reporting Server browser interface, you can encrypt passwords in configuration files, enable Secure Socket Layer (SSL) encryption for the TCP/HTTP listener, and encrypt data passed between the server and a remote server. For more information about encrypting communication between the server and the TIBCO WebFOCUS Client, see the *TIBCO WebFOCUS® Security and Administration Manual*.

Configuring Authentication

Authentication is the process of validating user credentials. Individuals who use WebFOCUS may have user IDs and passwords managed by other systems. WebFOCUS Managed Reporting and the Reporting Server can each be configured to authenticate users to an external security system, or to trust that authentication has already taken place. Users benefit because they do not have to sign in multiple times or manage separate user ID/password combinations. If a user signs in to WebFOCUS with credentials from an outside security package, the package provides some type of authentication confirmation to the WebFOCUS Client or the Reporting Server. This information may be in the form of a browser cookie or a logon ticket and may be needed by the Reporting Server in order for it to access and retrieve the data required by the WebFOCUS application.

Depending on your operating environment, user credentials can be validated by:

- The operating system on which the server is running (OPSYS security).
- A list of users stored in the server file named admin.cfg (PTH security).
- An LDAP or Active Directory repository (LDAP security).
- A DBMS (DBMS security).
- A custom user-supplied procedure (CUSTOM security).

Multiple security providers can be configured. You can configure additional providers using the Reporting Server browser interface Access Control page.

Configuring a New Security Provider

Security providers must be configured from the Reporting Server browser interface Access Control page.

To switch to security provider LDAP, DBMS, or CUSTOM, you first need to add a new provider under the appropriate security provider on the Access Control tree. When you add a new provider, the server updates the `edaserve.cfg` file with provider blocks named `LDAP_PROVIDER`, `DBMS_PROVIDER`, or `CUSTOM_PROVIDER`. These blocks are inserted with all attributes that apply to the provider type between `BEGIN` and `END` dividers.

Security provider PTH is configured during installation of the server and is the default security provider when a newly installed server is started. A server administrator can configure and activate other security providers.

The OPSYS security provider is always on the list of providers. In order to make it an active provider, the server must be authorized (privileged) to start with security OPSYS. There is no need for the additional step of creating a new provider as is needed for LDAP, DBMS, and CUSTOM providers. By default, the effective server administrator ID shown on the *Activate Providers* page is the server administrator ID. You can change this using the *Access Control Settings* page. Once the server is installed, you can start adding new PTH users and PTH groups, and you can assign PTH users to PTH groups.

Security provider CUSTOM supports using non-standard security storage. You can create DBMS tables to store users, groups, and passwords. You must then write procedures to read these tables and perform the basic security tasks. This security mechanism has to be created prior to configuring the CUSTOM security provider and must be available when the CUSTOM provider is added to the server provider configuration. It must perform the following standard security tasks:

- Retrieve the user ID and its password to authenticate the user on connecting to the server.
- Retrieve the list of users and groups on user and group registration.
- Retrieve the list of users in each group.
- Retrieve user membership in all available groups.

Once the provider is added, you can change security providers from the Access Control page. Right-click the *Security Providers* folder and select *Activate Providers* to switch to a newly added provider. At this point, the `edaserve.cfg` file is updated with the `security_provider` attribute that specifies the provider name.

Multiple security providers can be configured. When you configure OPSYS as one of multiple security providers:

- The server has to be authorized to run security OPSYS, and it needs to have a valid OS user ID and password for starting an agent.

- ❑ In a multi-provider security environment that includes an OPSYS security provider, other provider names should not match any domain name reachable by the OPSYS provider.

The Access Control *Settings* page allows you to enter valid OS credentials.

When you configure multiple providers, as shown in the following image, you choose one to be primary on the Access Control *Activate Providers* page. Drop-down lists are available for assigning the other providers as secondary providers or inactive. The primary provider defines the security under which the server is running:

Effective Server Administrator: PTH\svradmin			
Security Provider	Status	Accepts Trusted	Server Administrators
PTH<internal>	Primary	✓	PTH\svradmin, PTH\se02703
CUSTOM - cust01	Secondary	✓	
LDAP - ldap01	Secondary	✓	
OPSYS	Inactive		OPSYS\IBI\se02703

During connection to the server, all provider users are connected with a two-part user ID consisting of the provider name and the user ID, for example, MyLDAP\User1 or MyDBMS\User2, where MyLDAP or MyDBMS is the name of a configured secondary provider.

This naming convention is also used when users and groups are registered to server roles. Users and groups are registered with a two-part name. If the PTH provider is a secondary provider, PTH users and groups are registered (and used on connection) as PTH\userid and PTH\groupid.

Example: Security Provider Blocks in `edaserve.cfg`

The following is an example of a `security_provider` attribute and security provider block for starting the server with security LDAP (provider name MyLDAP):

```
security_provider = MyLDAP
LDAP_PROVIDER = MyLDAP
BEGIN
  ldap_host = ldaphost
  ldap_port = post
  ldap_secure_connection = n
  ldap_user_base = dc=ibi,dc=com
  ldap_user_scope = subtree
  ldap_user_class = person
  ldap_user_attribute = uid
  .
  .
  .
END
```

The following is an example of a security_provider attribute and security provider block for starting the server with security DBMS (provider name MyDBMS):

```
security_provider = MyDBMS
DBMS_PROVIDER = MyDBMS
BEGIN
    security_dbms = MSSQL
    security_connection = CON01
END
```

Configuring OPSYS Authentication

When security provider OPSYS is configured, the user credentials from the client connection are authenticated by the native security system of the operating system. The server then allocates a data access agent that impersonates the user so that access to files or other objects is controlled by the native system.

Profiles for operating-system users are supported on all platforms. On Windows, Active Directory groups are supported based on the win_primgroup_adsi setting.

Details about OPSYS security requirements are provided in the *TIBCO WebFOCUS® Reporting Server Installation* manual, where specific sections are provided for each platform. Refer to this manual for installation instructions and platform-specific setup steps. The following is a summary of the OPSYS requirements:

- ❑ **On UNIX**, the startup tscm300 executable must have root setuid authority.
- ❑ **On Windows**, the server must be started as a service under SYSTEM account (Administrators Group authority) or as a user ID with Windows Administrator authority.

On Windows, the Server Administrator password is required for this security provider. If the password is not provided in the configuration, the server starts in safe mode and displays a message to that effect. Multiple administrators are allowed. For more information on creating server administrators and other roles, see [Registering Users and Groups in a Role](#) on page 88.

- ❑ **On z/OS**, the MVS load library must be APF-authorized, and for HFS deployment, certain HFS executables must be given +a authorization.
- ❑ **On IBM i** (formerly known as i5/OS), the ownership of certain library files must be changed to have QSECOFR ownership and, therefore, authority.

Note that some system-specific settings in the edaserve.cfg file are provided to allow further adjustments of the authentication mechanism. Those relevant for some UNIX systems are:

- ❑ `require_logon_privilege` (AIX only. Allows no sign-in IDs to access the server.)

❑ `update_security_db` (AIX only. Registers failed sign-in attempts.)

For Windows systems, the `logon_method` setting (for interactive, network, or batch) is relevant to explicit connections.

Note: For more information on limiting user access to the Reporting Server browser interface, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

Preventing Unsecured Server Starts After Upgrades

If the environment variable `EDAEXTSEC` is explicitly set to `OPSYS`, and the server fails to start because it lacks system privileges to start configured for an `OPSYS` security provider, the server start aborts and error messages are written to the `edaprint` log file.

This feature prevents an unsecured server start after a software upgrade if any of the required post-upgrade reauthorization steps are missed on a UNIX, IBM i, or z/OS HFS deployment. This is not applicable to other platforms. The setting may be placed in any normal server start-up shell or profile that a site is using or in the server `edaenv.cfg` configuration file. The messages vary slightly by platform.

The `edaprint` messages are:

```
Configured security is 'ON' as set by EDAEXTSEC variable.
```

```
Server has no root privilege. (UNIX)
```

```
Server is not APF-authorized. (z/OS HFS)
```

```
TSCOM300.PGM has no QSECOFR authority. (IBM i)
```

```
Workspace initialization aborted.
```

```
(EDA13171) UNABLE TO START SERVER
```

Using Password Phrases for Authentication on z/OS

Starting with z/OS 1.8, IBM introduced password phrases which can be used for authentication, and starting with z/OS 1.10 password phrases can be used for TSO Logon as well. Password phrases can be from 14 to 100 characters long.

z/OS IDs can have both a password (1 to 8 bytes) and a passphrase (from 14 to 100 bytes). The two forms of credentials are independent. Changing one does not change the other. The server recognizes the type of credential being presented by its size (sizes between 9 and 13 are invalid on z/OS).

No special configuration is needed to enable password phrases.

When a user logs into the Reporting Server browser interface, there are two possible ways they can change their passwords or passphrases:

- ❑ On the initial Reporting Server browser interface Logon page. This option is disabled by default. For more information on enabling this option, see [How to Change Passwords on the Reporting Server Browser Interface Logon Screen](#) on page 48.
- ❑ After the initial logon, by choosing *Change Password* on the *My Console* menu. A new screen opens with three separated fields (old password, new password, and confirm new password).

Procedure: How to Change Passwords on the Reporting Server Browser Interface Logon Screen

By default, passwords cannot be changed on the Reporting Server browser interface Logon page because the following parameter is in effect:

```
password_change_wclogin = n
```

To enable the option to change passwords when logging into the Reporting Server browser interface:

1. Select *Access Control* from the Reporting Server browser interface menu bar.
2. Right-click *Access Control* on the tree and select *Settings* from the context menu.
3. Change the value of the `password_change_wclogin` parameter to *y*.

On the same configuration page, you can change the delimiter used to separate the user ID from the password or passphrase to any valid character accepted by the operating system. By default, the delimiter is a comma:

```
password_change_delimiter= ,
```

Set `password_change_delimiter = off` to disable this feature.

Once the option to change passwords on the Reporting Server browser interface Logon page is enabled, users can change their passwords or passphrases by entering the following in the password field of the Reporting Server browser interface Logon page (if the password delimiter has been changed, use that character between the old and new passwords instead of the comma):

```
old_password,new_password
```

When `password_change_wclogin = y`, passwords and passphrases cannot contain the password delimiter character.

When `password_change_wclogin = n` (the default), the setting for `password_change_delimiter` is ignored, and passwords can contain all valid characters allowed by the operating system.

Configuring Linux and AIX Pluggable Authentication Modules (PAMs)

In a Linux or AIX environment, OPSYS security can be configured to use Pluggable Authentication Modules.

1. Go to the *Access Control* page.
2. Right-click the OPSYS security provider and select *Properties* from the context menu.
3. Select `y` from the `security_pam` drop-down list.

This parameter is only applicable for the Linux and AIX operating systems. The parameter defines if the server uses the Pluggable Authentication Modules mechanism to implement security. If `y`, the server uses PAM calls, if `n` the server uses native UNIX security calls.

4. Click *Save and Restart Server*.

Using Kerberos for Single Sign On

A server started with security provider OPSYS can be configured for Kerberos connections.

Procedure: How to Implement Single Sign On Kerberos Security on Linux

1. On the *Access Control* page, right-click the OPSYS provider, and select *Properties* from the context menu.

The OPSYS Security Configuration page opens.

2. In the `krb5_srv_principal *` field, enter your server principal used for Kerberos security.
3. Click *Save*.

The `edaserve.cfg` file is updated with this attribute.

4. On the *Workspace* page, expand the *Special Services and Listeners* folder, right-click *TCP/HTTP*, and select *Properties of HTTP*.
5. Check `SECURITY=KERBEROS`.
6. Click *Save and Restart Listener*.

As a result of this configuration step, the `odin.cfg` file located in `EDACONF/etc` is updated for the `HTTP_LST` node.

```
;HTTP Listener
NODE = LST_HTTP
BEGIN
PROTOCOL = HTTP
PORT = 8101
CLASS = AGENT
SECURITY = KERBEROS
END
```

When the server is started, a user can connect to the Reporting Server browser interface from Internet Explorer without a prompt for user ID and password. The *Login Info* shows connection type *Kerberos*. The connection is done using the Kerberos ticket from the browser. The connected user ID is derived from this ticket.

If the Linux operating system uses the System Security Daemon Service (SSDS), there is no need to maintain a local OPSYS user ID with the same name as the Kerberos user ID. However, if the Linux operating system does not use SSDS, the connection to the server requires that there is a local OPSYS user ID with the same name as the Kerberos user ID on the operating system running the server. This user ID is used for `tscom3` process impersonation.

If a user signs off from the Kerberos connection, the user can make explicit connections with the local Unix user ID and password. Connection with another Kerberos user ID as an explicit connection will not work.

Procedure: How to Implement Single Sign On Kerberos Security on Windows

1. Configure the Reporting Server to start with security OPSYS.

The server must be started as a service that is set to run under the LocalSystem account.

2. On the Workspace page, expand the Special Services and Listeners folder, right-click *TCP/HTTP*, and select *Properties of HTTP*.
3. Expand the Security section and click the *Kerberos or NTLM* check box for SECURITY.

As a result of this configuration step, the `odin.cfg` file located in `EDACONF/etc` is updated for the `HTTP_LST` node.

```
HTTP Listener
NODE = LST_HTTP
BEGIN
PROTOCOL = HTTP
PORT = 8117
CLASS = AGENT
SECURITY = KERBEROS
END
```

The connection from the browser will be done using Kerberos or NTLM, depending on whether a Kerberos connection can be made.

Configuring PTH Authentication

By default, the server installation process establishes PTH security with a PTH server administrator user ID and password. If you deleted the PTH server administrator ID, you can reconfigure PTH security.

The PTH security provider only controls access to the Reporting Server browser interface. When the server is configured for this security provider, there is no impersonation by data agents or authentication of a non-Reporting Server browser interface connected user. From the operating system point of view, all server processes run as a single-user ID, and access to the Reporting Server browser interface is controlled by authenticating names against those defined in the `admin.cfg` file, unless `authenticate_all_pthuser` is set to `y`. The setting is available on the PTH Security Configuration page, which can be accessed by right-clicking *PTH* in the *Security Providers* folder on the Access Control resources tree.

You must configure the Server Administrator password *before* starting a server configured for a PTH security provider, either by providing it at installation time or by configuring it in the Reporting Server browser interface. For details, see [Registering Users and Groups in a Role](#) on page 88.

The PTH security provider supports profiles for users and groups on all platforms.

Note: For more information on limiting user access to the Reporting Server browser interface, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

Procedure: How to Configure Password Properties for a PTH Security Provider

When the Server Administrator creates or updates a user under a PTH security provider, by default the password never expires, can be any length, and is not case-sensitive.

The Server Administrator can configure a new or existing user password to expire after a specified number of days, can set a minimum length for passwords, can enable case-sensitive passwords, and set other password properties.

1. To assign an expiring password, click *Access Control*.
 - a. Right-click *PTH* and click *Manage Users/Groups*.

The *PTH Users and Groups Management* page opens.
 - b. If the user does not yet exist, create the user. Then, click *Properties* on the menu bar.

By default, *Password never expires* is checked.

- c. Uncheck Password never expires.

 Password never expires' and '? Account disabled'. At the bottom right are 'Cancel' and 'OK' buttons."/>

PTH<internal> User ✕

PTH<internal> User? PTH\user1

Description?

E-Mail?

Password? (required) expires in 42 day(s)

confirm password? (required)

? Password never expires

? Account disabled

When the password expires, the user gets a *password expired* message on the logon screen and is provided with a *New Password* field in order to enter a new password.

At this time, the new password and new password change date are recorded in the `admin.cfg` file.

Note that changing the password by selecting *Change Password* from the *My Console* menu also resets the password change date.

- d. Click *OK*.
 - The current date is saved in the `admin.cfg` file as `admin_passdate` for this user. This is the last password change date.
 - e. Click *Save*.
 - f. Optionally, configure a warning message to begin displaying a specified number of days before the password expires, as described in [Sign-in Password Expiration Warning](#) on page 61.
2. To configure other password properties:
- a. Right-click the PTH Security Provider and click *Properties*.

The Security Configuration PTH page opens

authenticate_all_pthuser	?
<input type="text" value="y"/>	▼
pthpass_casesensitive	?
<input type="text" value="n"/>	▼
pthpass_mixedchars	?
<input type="text" value="n"/>	▼
pthpass_lifetime	?
<input type="text" value="60"/>	
days	
pthpass_lifetime	?
<input type="text" value="0"/>	
characters	
pthpass_unique	?
<input type="text" value="0"/>	
# passwords	
pthpass_lockout	?
<input type="text" value="0"/>	
failures	
pthpass_lockout_duration	?
<input type="text" value="0"/>	
minutes	
trust_ext	?

- b. Set the following properties.
- authenticate_all_ptuser.** Valid values are *y* and *n*. Defines whether the server will authenticate all connections, including Basic Users, to the Reporting Server browser interface. The default value is *y*.
 - pthpass_casesensitive.** Valid values are *y* and *n*. Defines whether authentication of passwords will be case-sensitive. The default value is *n*.
 - pthpass_mixedchars.** Valid values are *y* and *n*. Defines what characters should be required in a password. If the value is *y*, a password must contain at least one special character, one numeric character, and one letter (one uppercase and one lowercase letter, if `pthpass_casesensitive` is set to *y*). The default value is *n*.
 - pthpass_lifetime.** Defines the maximum password lifetime in days. Zero (0) means no expiration. The default is 60.
 - pthpass_length.** Defines the minimum number of characters that must be present in a password. Zero (0) means no minimum requirement. Zero is the default value.
 - pthpass_unique.** Defines the number of password cycles during which a password cannot be reused. The value zero (0) means passwords can be reused. Zero is the default value.
 - pthpass_lockout.** Defines the number of failed sign-in attempts that will cause a user account to be locked. The value zero (0) means no limit. Zero is the default value.
 - pthpass_lockout_duration.** Defines the number of minutes that a locked-out account remains locked out before automatically becoming unlocked. The value zero (0) means that the account will be locked out until an administrator explicitly unlocks it. Zero is the default value.
 - trust_ext.** Valid values are *y* and *n*. Specifies whether the server should accept trusted client connections. If there are multiple security providers, some may allow trusted connections and some may not. In these cases, if the trusted connection is made using a provider without the ability to accept trusted connections, the user will get an authentication error. The default is *y* for security OPSYS (except Windows security OPSYS), *n* for Windows security OPSYS and all other security providers.
- c. Click Save.

Disabling a PTH Security Provider Account

You can disable an account that is registered in the admin.cfg file under the PTH security provider in order to prevent a user from signing in to the server with that account.

To disable a PTH account:

1. Right-click *PTH* under the *Security Providers* folder, and select *Manage Users/Groups* from the context menu.

The PTH Users and Groups Management page opens, listing all registered users on the left, and all registered groups on the right.

2. Select a user and click *Properties*.

The PTH User properties dialog box opens.

3. Click *Account disabled*, as shown in the following image.

The screenshot shows a dialog box titled "PTH<internal> User" with a close button (X) in the top right corner. The dialog contains the following fields and options:

- Field: PTH<internal> User? PTH\svaa1
- Field: Description? (empty text box)
- Field: E-Mail? (empty text box)
- Field: Password? (password masked with dots) (required) expired
- Field: confirm password? (password masked with dots) (required)
- Option: ? Password never expires
- Option: ? Account disabled
- Buttons: Cancel, OK

4. Click *OK*.

An x is displayed in the Disabled column for that user on the PTH Users and Groups Management page.

5. Click *Save*.

Configuring DBMS Authentication

With a DBMS security provider, there is no impersonation by data agents, but connection credentials are authenticated against a configured DBMS. This technique is called *password passthru*, as user IDs and passwords supplied by the client are passed to the DBMS for authentication.

A DBMS security provider supports profiles for users on all platforms. Profiles for groups are not supported.

Note: For more information on limiting user access to the Reporting Server browser interface, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

Configuring LDAP Authentication

With an LDAP (Lightweight Directory Access Protocol) security provider, there is no impersonation by data agents, but connection credentials are authenticated against the established directory services. From the operating system point of view, all server processes run as a single-user ID.

An LDAP security provider is presently supported on Windows, UNIX, IBM i, and the Unified Server (HFS and PDS deployments).

An LDAP security provider supports profiles for LDAP users and LDAP groups.

Note: For more information on limiting user access to the Reporting Server browser interface, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

Configuring Custom Authentication

CUSTOM security providers are used when the installation maintains non-standard repositories for users and groups, that is, when the standard methods like OPSYS and LDAP do not apply. An example of such a setup can be having RDBMS tables containing valid user IDs and their encrypted passwords, and also containing the mapping of users to groups. The customer may choose to access these tables using SQL SELECT statements, SQL Stored Procedures, or Web Services. Typically these methods already exist and are used by other non-TIBCO components.

To enable the custom server security provider, the administrator needs to provide code that allows the server security module to perform these tasks:

- Authenticate the user based on user ID and password.
- Obtain all groups for a user (used at run time).

- ❑ Obtain all groups in the system (used for administrative actions).
- ❑ Obtain all users for a group and all users in the system (used for administrative actions).

The code for performing these tasks is written in the WebFOCUS language as TABLE FILE or SQL SELECT commands against the synonyms representing the SQL table, SQL Stored Procedure, or Web Service. The code can be debugged by running the server with security OFF and running the glue code from application folders. Once debugged, the code is deployed in EDACONF/catalog/custom, and the CUSTOM provider is made active.

A tutorial sample called *WebFOCUS - Custom SQL Security Provider* is available that creates examples of user and group storage and custom procedures that can be used as prototypes. To generate a tutorial, right-click an application on the Application page, point to *New*, and click *Tutorials* from the context menu.

Encoding Passwords for a Custom Security Provider

You can configure Custom Provider authentication to encode passwords so that they cannot be viewed when they are passed as parameters to the authentication procedure.

1. On the Access Control page, right-click a Custom Security Provider and select *Properties* from the context menu, or right-click the CUSTOM folder and select *New* from the context menu.

The CUSTOM Security Provider Configuration page opens.

2. Select *y* from the *cust_hashpasswd* drop-down list.
3. Click *Save*.

When a user is authenticated, the server will send SHA-256 hashed passwords to the authentication procedure. This means that the passwords will be transformed into values that will be unusable for signing in to the Reporting Server browser interface and server. The hashed values have to be stored as passwords in the SQL database being used for authentication. As an alternative, you can create a utility called from the CUSTOM provider authentication procedure that decodes hashed passwords before sending a request to the SQL database for authentication.

Authenticating Users Across Multiple Security Providers

The Reporting Server can search across multiple LDAP sources, DBMS providers, an OPSYS provider, and a PTH server when authenticating users.

When authenticating or assigning privileges for a provider that is not the primary provider, the user ID is a two-part name consisting of the provider name and the user ID:

provider\userid

The authentication is done based on a two-part name.

The Server Administrator can add and remove security providers from the list at any time.

For instructions on configuring PTH and LDAP security providers, see [Configuring PTH Authentication](#) on page 51 or [Configuring LDAP Authentication](#) on page 56.

Enabling Trusted Connections

You can enable trusted connections separately for each security provider.

You can configure a security provider to accept trusted connections by setting its `trust_ext` value to `y`.

1. Right-click a security provider on the Access Control tree and click *Properties* on the context menu.

The Configuration page for that security provider opens.

2. From the `trust_ext` drop-down list, select `y`.

This setting is available for all types of security providers.

3. Click *Save and Restart Server*.

Identifying a Default Provider for Trusted Groups

When multiple security providers are configured, one must be identified as the default provider for trusted connections when WebFOCUS or other client software sends a trusted group ID without a security provider to the server.

To set the default provider for trusted groups:

1. On the Access Control page, click *Settings*, then *Access Control* on the ribbon, or right-click the *Access Control* folder, point to *Settings*, and select *Access Control* from the context menu.

The Access Control Settings page opens, as shown in the following image.

General	
Effective Server Administrator?	PTH\srvadmin
trusted_group_default_provider?	PTH (accepts trusted)
group_profile?	primary
default_admin_role?	Basic User
prepend_provider_name?	y

2. You can select a security provider that accepts trusted connections from the `trusted_group_default_provider` drop-down list. You must select a provider that is configured to accept trusted connections. You can also enter IP addresses to be mapped to a specific provider, if no provider name is supplied by the client.

Setting an Anonymous User ID

An authorized server administrator can set an anonymous user ID that can access the Reporting Server browser interface when the user ID/password fields on the Reporting Server browser interface sign-in screen are blank and the user clicks *Sign in*.

This anonymous user ID provides the ID and, in turn, the user rights for the Reporting Server browser interface and the server. Further configuration for the anonymous user can be achieved by creating a user profile for the anonymous user ID. If an anonymous ID matches a user in the `admin.cfg` list, the applicable role and privileges are applied.

Note: On Windows, for security provider OPSYS only, you must turn IWA security off to use this feature. From the *Workspace* page, open the *Special Services and Listeners* folder, right-click *TCP/HTTP*, and select *Properties*. Click the *Security IWA* check box to deselect it.

1. Log on to the Reporting Server browser interface using a user ID that has server administration rights.
2. Go to the *Access Control* page.
3. Click *Settings*, then *Access Control* on the ribbon.
The Access Control Settings page opens.
4. Enter a user ID in the *anonymous_id* field.

5. Enter a password in the *anonymous_pass* field and repeat it in the *confirm password* field.
6. Click *Apply and Restart Server*.

Configuring User Password Settings

You can configure the following password settings for any security provider:

- View or change the password separator character.
- Enable users to change passwords from the Reporting Server browser interface sign-in page.
- Control the start time of warning messages on operating systems that support password expiration.

Procedure: How to View or Change the Password Separator Character

When an application has the option to change a user password, the old and new passwords are sent to the server using a separator character, by default, a comma (.). For example:

old_password,new_password

The separator character is defined by the *password_change_delimiter* field. The separator needs to be reset if the current separator character is contained in the password itself or if it is allowed by the server.

This feature is not supported when Security is OFF.

1. Go to the *Access Control* page.
2. Click *Settings*, then *Access Control*.

The Access Control Settings page opens.

3. In the *password_change_delimiter* field, enter a single character to use as the delimiter between an old and new password. (The default character is a comma.)

You cannot use the designated *password_change_delimiter* character in a password.

Note: If you enter *off* in the *password_change_delimiter* field, users will not be able to change their passwords on the Reporting Server browser interface login page, but all characters will be supported in passwords.

4. Click *Apply and Restart Server*.

Procedure: How to Enable Password Changes From the Reporting Server Browser Interface Sign-in Page

A Server Administrator can allow users to change their password from the Reporting Server browser interface sign-in page. By default, users cannot change their passwords from the Reporting Server browser interface sign-in page.

1. Sign in to the Reporting Server browser interface with Server Administrator privileges.
2. Go to the *Access Control* page.
3. Click *Settings*, then *Access Control*.

The Access Control Settings page opens.

4. Select *y* in the *password_change_wclogin* field. The default is *n*.
5. Click the *Apply and Restart Server* button.

Sign-in Password Expiration Warning

On Operating Systems that support password expiration, you can display a warning message a specific number of days before the expiration date. The message will appear after the initial sign-in screen. You must then click the *Continue* button to open the Reporting Server browser interface. Users can employ standard tools to change their passwords before expiration occurs.

Note: The password expiration warning is supported only on operating systems where IDs are configured to expire, and this extended security feature is active for the user ID. However, expiration warning is not currently supported on Windows platforms.

How to Initiate Notification of an Imminent Password Expiration

1. Go to the *Workspace* page.
2. In the resources pane, expand the *Special Services and Listeners* folder, then right-click *TCP/HTTP* and select *Properties or HTTP* from the context menu. The Listener Configuration pane opens.
3. Expand the *Sessions Control* group.
4. In the *PASS_EXPIRE_NOTIFICATION* field, enter the number of days prior to expiration that you want the warning to begin to appear when a user whose password is about to expire logs in.
5. Click *Save and Restart Server*.

Configuring Privileges and Other Authorizations

WebFOCUS users are authenticated against one of the supported authentication sources (OPSYS, LDAP, DBMS, PTH internal, or CUSTOM). The underlying security system may divide users into groups. For example, for OPSYS security the groups are defined in the operating system, and in LDAP they are defined in the LDAP database. With a PTH security provider, the server file named `admin.cfg` defines groups. A CUSTOM provider executes custom requests against user-defined security storage, for example, an SQL database.

To assign privileges to the server application and data files, you must register groups of users to a server role and customize the privileges. In exceptional cases, you may register individual users to a server role and customize its privileges. The authorizations then depend on the privileges assigned to the role.

With any type of security, access to the following categories of resources must be controlled:

- ❑ **DBMS data.** Access is controlled by the CONNECTION attributes in effect at the time the data is accessed. The connections are stored in user or group profiles or in the server `edasprof` profiles.
- ❑ **Access to Reporting Server browser interface system functions.** This category includes functions such as change of server profiles, creation of metadata, or starting the servers. Access to these functions is defined by the role assigned to the user or group and is stored in the `admin.cfg` file.
- ❑ **WebFOCUS repository files that constitute the application.** These files include FOCEXECs, Master and Access Files, HTMLFORMs, and application files that are stored in the application directories.

The Server provides an *Access Control* function that enables the Server Administrator to create and edit permissions for different categories of users and groups. These permissions define whether specific users and groups can read, write, and/or execute WebFOCUS repository files and perform various Reporting Server browser interface system functions.

In addition to users and groups, the server supports security subjects called *roles*. The server comes with a set of standard roles, and these roles come with a fixed set of general privileges. For more information about these standard roles, see [Configuring Roles](#) on page 85.

The Server Administrator can register roles, groups, and users. When registered, they are assigned a set of general privileges and file access privileges. By default, groups inherit privileges from the roles they are assigned, and users inherit their privileges from the groups or roles to which they are assigned. The Server Administrator can create new roles and customize privileges for any role, group, or user. The Server Administrator always has full privileges, and those privileges are not adjustable.

It is considered good practice for a large user community to define access rules on the group level and not on the individual user level. This technique allows volume user provisioning to be handled outside of the WebFOCUS Reporting Server software (for example, in the LDAP server). The WebFOCUS Reporting Server in general registers group access rights, plus some individual users as an exception, such as managers or special project people.

The following defines the order of precedence for access rights:

- ❑ User access overrides group access, and group access overrides role access.
- ❑ If the user belongs to multiple groups, the server `group_profile` setting determines whether only the primary group is in effect or whether the user can select a group at logon.

Configuring DBMS Authorization

With DBMS authorization, the main object to be protected is the DBMS resident data that the server reads and writes. Access control is implemented using ENGINE CONNECTION_ATTRIBUTES commands that define the security attributes that are used by the server agent for connecting to the DBMS. The connection security type depends on the DBMS and can be one of three subtypes:

- ❑ **Explicit.** The userid/password is stored in a profile (encrypted).
- ❑ **Password passthru.** The userid/password combination or cookie, such as MYSAPSS02, that were provided during the connection to the server is passed to the DBMS connection.
- ❑ **Trusted.** This type of connection applies to security provider OPSYS where the agent operating system process impersonates the user, and the DBMS connection is derived from that process.

The CONNECTION_ATTRIBUTES command itself can be defined in the user, group, role, or server profile, and they override each other in this order.

It is common initially to create connections in the edasprof server profile. They will then be inherited by all users. Subsequent configuration steps may define connections on other levels, such as the group or user level.

The CONNECTION_ATTRIBUTES command in effect passes the credentials to the DBMS, and the DBMS server ensures the correct access control: read and read/write rules on the DBMS tables, views, columns, and rows.

For more information about the CONNECTION_ATTRIBUTES command, see the *TIBCO WebFOCUS® Adapter Administration* manual.

Configuring DBA Security

The server engine can add an additional level of access control by defining DBA rules in synonyms to restrict access to columns and certain data values. This applies only to data access using the TIBCO WebFOCUS® language, and not to Direct SQL Passthru requests. For more information about Server DBA, see the *Describing Data With TIBCO WebFOCUS® Language* manual.

Calculating Privileges for Any Registered or Unregistered User or Group

The server administrator can calculate the privileges of any registered or un-registered user or group on the server by right-clicking the Access Control tree and selecting *General Privileges* or *Directory/File Privileges* from the context menu or by clicking *General Privileges* or *Directory/File Privileges* on the ribbon.

The *Show Privileges* page opens.

You can select a security provider and whether to show the privileges of a user, group, or role under that security provider. You must enter a valid ID for the user, group, or role.

If you enter an unregistered user or group, then the server will calculate the privileges based on user membership and its group registration.

When you click *Next*, the server calculates the appropriate privileges and returns a page of Properties. This page also indicates whether the security subject you chose is registered and, if not, which privileges it inherits. If your server is configured with `profile_setting=all`, and you select a user who belongs to multiple registered groups, the privileges of all will be combined in the display.

Permissions for Server Application Files and Directories Using a Non-OPSYS Security Provider

The Server Administrator can specify read/write, execute, and list privileges for files in application directories and for the directories themselves. This control is in addition to operating system file access control which follows the current umask setting.

The Server Administrator can view or edit File, Directory, and General privileges by clicking the *Directory/File Privileges* or *General Privileges* button on the ribbon, or by right-clicking a role, group, or user on the Access Control page and selecting *Directory/File Privileges* or *General Privileges*.

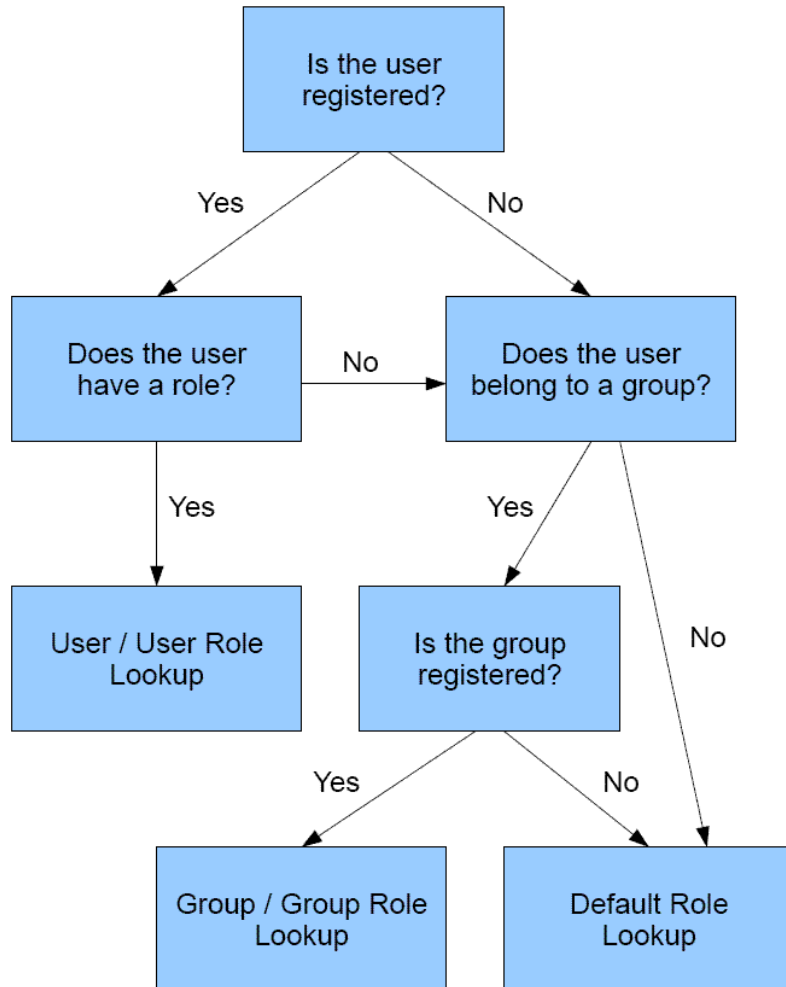
Non-Server-Administrator users can right-click any object on the left panel and select *Properties* to see their privileges for that file or directory object. The Server Administrator has the option to manage privileges by right-clicking an object.

By default, for compatibility with previous releases, file access control allows complete access to all files and directories. However, the Server Administrator can reset this default by right-clicking a row in the table and clicking *General Privileges* or *Directory/File Privileges*. For information, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

After initial server installation, the Server Administrator should review individual privileges for each role and customize them if needed. For example, the write privileges for basic users to applications on the APPPATH can be removed.

File privileges can be defined on the lowest level, such as a specific file and specific user, but the recommended policy is to define them on the application folder level and on the role and group levels to reduce administrative overhead. Privileges defined on the application level will apply to (be inherited by) all files in this application and by all nested (lower level) application directories, unless overridden on a lower level. The same applies to roles, groups, and users. Privileges defined on the group level will be inherited by the users in the group, unless overridden by explicit user registration.

The following schematic shows the flow of privilege assignments and inheritance.



In effect, when the system opens a file such as *app1/app2/car.mas* for user *user1* in *group1* belonging to *role1*, the server uses a specific search path until the privileges are found. The server administrator can also see the directory and file privileges from the Application Directories page by right-clicking an application and selecting the *Privileges* option.

Users without administrator rights can see the privileges from the file or directory *Properties* option.

Example: Search for Privileges for Users Who Are Not Registered

The following describes the search path if user1 is not registered or has not been assigned a role in the admin.cfg file. Assume the user is in group1 and group1 is registered under role1:

1. App1/app2/car.mas for group1.
2. App1/app2 for group1.
3. App1 for group1.
4. * for group1
5. App1/app2/car.mas for role1.
6. App1/app2 for role1.
7. App1 for role1.
8. * for role1.

If group1 were not registered, the user would be assigned the default role.

Example: Search for Privileges for Registered Users

If user1 is registered in the admin.cfg file, the server uses the following search path. The lines ending with an asterisk (*) denote a recommended practice:

1. App1/app2/car.mas for user1.
2. App1/app2 for user1.
3. App1 for user1.
4. App1/app2/car.mas for role1.
5. App1/app2 for role1.
6. App1 for role1.
7. * for role1.

If the user is registered under a role, the server does not consider group privileges when access to the file or directory is calculated.

Reference: Understanding Access Control Permissions

When file and directory access control privileges are being set, the registration of physical file locations is not saved using a full path reference in all cases. When a location is in a folder under one of the internal server locations EDACONF, EDAPRUFU, APPROOT, or EDAPRUFU, it is registered as a reference that is relative to that server location. The physical paths to EDACONF, EDAPRUFU, EDAPRUFU and APPROOT are defined when the server is installed.

For example, if the privilege is for the directory D:\ibi\svr90\wfsV8\etc, and if the EDAPRUFU location has been established as D:\ibi\svr90\wfsV8, any privilege registered for that location will be registered using a relative reference based on EDAPRUFU:

```
admin_privilege = (EDAPRUFU)\etc;PRRUN
```

Then, if a new server release is installed in a different file system, the privilege will be applied to the correct directory in the new file system because the EDAPRUFU location will contain the correct directory reference for that system.

Registration of files and directories located outside of EDACONF, EDAPRUFU, APPROOT, and EDAPRUFU locations are registered as full physical paths.

Access Control parameters are saved in the server admin.cfg file.

```
{admin_id|admin_group} = name  
admin_privilege = object; privilege_name, privilege_name[, ...]
```

where:

name

Is a user (admin_id) or group (admin_group) name.

object

Defines the directory or file path.

Can be one of the following:

- ❑ Physical path to the file or directory.

For example:

```
c:\myapp\app1 c:\myapp\myfile.foc
```

where * is a token that designates all physical files on the system.

- ❑ A path relative to one of the server internal locations EDACONF, EDAPRUFU, APPROOT, or EDAPRUFU.

For example:

```
(APPROOT)\app1 (APPROOT)\ibisamp\car.foc
```


privilege_name

Defines the type of access. Can be one of the following:

AREAD. Reads and displays the content of a file to the user. For example, the user can see (but not change) procedures and synonyms. The user needs read permission to open data files.

ARWRT. Allows the user to read, edit, and write procedures and metadata.

PRRUN Execute. Typically given to end users who need to execute FOCEXECs and use Master Files and relevant utility procedures. However, the data file used by a FOCEXEC and Master File only requires read permission.

ALIST. Allows the user to list and view files and folders on the application tree or the output of a deferred execution run.

ANONE. Revokes all permissions.

Note that definitions in the `admin_privilege` strings can contain relative path references for application directories under `APPROOT` and other internal locations. However, fully-qualified physical paths are also respected.

Example: Useful Combinations of Permissions

1. Full Access - read, write, execute, list

This setting is typically used by developers on applications that they develop and create directly.

This is the default for all files not explicitly customized, unless the default is reset. To explicitly specify full access, issue the following setting in the `admin.cfg` file:

```
admin_privilege = directory/file_path; ARWRT, AREAD, PRRUN, ALIST
```

2. Read, execute, and list.

This permission can be used by developers for other utility application directories that they invoke while developing their own.

For example, a developer can keep common synonyms and utility FOCEXECs in such a directory. Of course, the person that creates such a directory needs full access to it, but all others only need read and execute. Read is needed to open data files.

```
admin_privilege = directory/file_path; AREAD, PRRUN, ALIST
```

3. Execute.

This permission is typically given to end users who need to execute FOCEXECs and use Master Files and relevant utility procedures.

```
admin_privilege = directory/file_path; PRRUN, ALIST
```

4. Execute and hidden.

This is the same as Execute except that it is not shown on the Reporting Server browser interface.

```
admin_privilege = directory/file_path; PRRUN
```

5. List and View.

This allows the user to view reports from deferred jobs. This privilege is lower than execute.

```
admin_privilege = directory/file_path; ALIST
```

6. No Permissions.

This revokes all permissions from the directory or file.

```
admin_privilege = directory/file_path; ANONE
```

Example: Sample Permissions

The privileges for locations relative to EDACONF, EDAHOME, APPROOT, and EDAPRFU are set using the following syntax:

```
admin_privilege=(APPROOT)/baseapp;AREAD, ARWRT, PRRUN, ALIST
```

However, existing privileges set using the full physical name are respected.

These examples use the following directories and files:

```
/ul/pgmabb/myapp  
/ul/pgmabb/myapp/data1.dat
```

The following shows a sample permission distribution list:

```
admin_id = chief_developer
BEGIN
admin_level = APP
admin_privilege = *; ARWRT,AREAD,PRRUN,ALIST
admin_privilege = (APPROOT)/work_app; AREAD,PRRUN,ALIST
END
admin_group = team_developer
BEGIN
admin_level = APP
admin_privilege = (APPROOT)/main_app; AREAD,PRRUN,ALIST
admin_privilege = (APPROOT)/synonym_app ; AREAD,PRRUN,ALIST
admin_privilege = (APPROOT)/work_app; ARWRT,AREAD,PRRUN,ALIST
admin_privilege = /u1/admin/server/prod_data ; AREAD,ALIST
END
admin_group = end_user
BEGIN
admin_level = USR
admin_privilege = (APPROOT)/main_app; PRRUN,ALIST
admin_privilege = (APPROOT)/ synonym_app ; PRRUN
END
```

These security declarations define the following roles and permissions:

- A user with the ID chief_developer has full access to all applications, including main_app and synonym_app, but has no write access to work_app for the team_developer group.
- A user that belongs to the group team_developer develops applications in their own directory, but needs Read and Execute access to main applications to see what was done before. A user with this role has read-only access to production data in /u1/admin/server/prod_data. When the developer user is finished, the chief_developer copies the procedures to the main_app.
- The end_user group contains end users who need to execute the procedures from the main_app, but do not need to see synonym_app on the Reporting Server browser interface.

Reference: Access Control vs. APP PATH

APP PATH defines two main aspects of the system:

- Search path for applications at run time for one-part file reference, such as TABLE FILE CAR.
- List of applications shown on the Reporting Server browser interface and Data Migrator desktop interface application tree for the Application page.

The privileges are on the physical directories, not the application, so if the application is mapped to a different directory, the privileges from the new directory will be in effect.

Note that with Access Control, all such access is subject to Access Control rules which take precedence over APP PATH:

- ❑ A one-part file reference that searches APP PATH and finds a file in a directory without access will skip it and search further in APP PATH.
- ❑ A two-part file reference or FILEDEF attempting to access a file or directory without permission will produce an error.
- ❑ The Reporting Server browser interface and Data Migrator desktop interface will gray out actions such as saving a procedure or synonym where they are prohibited.

For example, user u1 has the following in their profile where both app1 and app2 contain the xxx.mas file and the f1.fex file:

```
APP PATH app1 app2 app3,
```

In the admin.cfg file for user u1, access to these applications are defined as:

```
admin_privilege = (APPROOT)/app1; AREAD, PRRUN
admin_privilege = (APPROOT)/app2; PRRUN
admin_privilege = (APPROOT)/app3; ARWRT, AREAD, PRRUN
```

When user u1 tries to execute the FOCEXEC f1 (EX f1), the one from app1 will be executed.

When user u1 tries to issue TABLE FILE xxx, the app1/xxx.mas will be opened, since it only needs Execute privileges.

When user u1 issues the following command, the app1/xxx.mas will be opened, since it requires Read privileges:

```
EDAGET MASTER, xxx, T, ANY
```

When user u1 tries to issue the following command ENCRYPT FILE XXX, the app1/xxx.mas will be read, and an attempt to rewrite it in place will be made. Since user u1 does not have sufficient privileges, an error will occur and an error message will be issued.

Protection against hacking:

- ❑ The admin.cfg file itself cannot be accessed with a WebFOCUS command. This rule is enforced in the code.
- ❑ All files that are not covered by explicitly registered privileges are covered by the system-wide default.

Reference: Permissions for FOCCACHE and EDATEMP

EDATEMP is a location where the server creates temporary files during agent connection. All of these files are deleted on disconnect to provide a clean environment for the next connection.

FOCCACHE stores temporary files during execution of one Web session. These temporary files can be stored in FOCCACHE for multiple agent connections, and data can be accessed through the Reporting Server browser interface and WebFOCUS Reporting Server. FOCCACHE is an internally mapped application directory to a unique physical location for each Web session.

Both of these locations have a full set of permissions for a connected user unless the default location has been overridden by HOLDDATA, HOLDMETA, and remapping of FOCCACHE from its default location using APP commands. In this case, the access controlled permissions are applied based on the admin role, group, or user.

Reference: Access Control Implications For Scheduling

The TIBCO® Data Migrator scheduler initiates requests based on procedures (flows) with scheduling headers that it finds in the application path. The procedures are submitted for execution using the user ID and related profiles that depend on the value of the sched_run_id keyword.

The values are set using the Scheduler Manage Tenants page. For information, see [How to Manage Tenants](#) on page 739.

Depending on the entries configured on the Manage Tenants page, syntax similar to the following is created in the edaserve.cfg file.

```

sched_scan_id = PTH\srvadmin
BEGIN
  sched_run_id_type = user
END
sched_scan_id = ldap01\user1
BEGIN
  sched_run_id_type = server_admin_id
END

```

The encrypted password is taken from the admin.cfg file. If the user is registered without a password, trusted connection is used for the LDAP and OPSYS security modes on all operating systems except Windows.

If the installation has a large number of users, it is good practice to protect which procedures the users are allowed to schedule. The method to accomplish this is to define a special application directory, for example, *schedule_app_1*, that contains process flows with scheduling headers only. The actual procedures can be in any other application directory including a home directory, for example, *app2*. The process flows will contain references to the actual procedures, for example:

EX *app2/myrpc*

A special scheduling directory such as *schedule_app_1* can be protected with file access privileges so that only a designated user ID, for example *user1* can write to it and create and modify schedules in it. The *user1.prf* profile should have *app1* in its application path.

The scheduler is configured to scan this application directory by setting the *sched_scan_id* to *user1* on the Manage Tenants page.

This ensures that only the schedules located in *schedule_app_1* are in effect and that *user1* is the only user controlling the scope of scheduling.

File Permissions for an OPSYS Security Provider

Impersonation is when the server executes code in the context of an authenticated and authorized user. Access to files is controlled by the operating system security features in the native file systems. Files are protected according to the permissions set for the impersonated user.

Users who have server privileges to make configuration changes (update configuration files such as *edaserve.cfg*, *admin.cfg*, *odin.cfg*, or *edasprof.prf*) must also have operating system privileges to update those files.

The main file areas within a configuration are organized into four locations:

- EDAHOME (the software installation structure).
- EDACONF (the specific software configuration structure).
- APPROOT (the file structure for holding application-related programming files).
- EDATEMP (the run-time user area for server processes, which is normally located under the EDACONF area).

Permissions for new files depend on how the server is installed and started.

- On UNIX, z/OS (HFS deployment), and IBM i**, the permission mask setting of the process (*umask*) that started the server is inherited down to the creation of each subdirectory or file with a few exceptions, and thus defines the permissions used for object creation.

- ❑ **On z/OS (PDS deployment)**, the rules (permissions) are set by the security package (RACF, ACF/2, or Top Secret).
- ❑ **On Windows**, the permissions are inherited down from directory to subdirectory, and the permissions for any object created are taken from its parent directory. So unless specific permissions are applied at any level, default permissions will depend on where the server was installed.

Reference: **Consequences for Files Under the EDAPHOME/EDACONF Hierarchy**

All users who are expected to use the server must have Read/Execute access to files under the hierarchies. Server administrators that share common resources (such as admin.cfg), should be members of a group that provides writable access to those resources. Non-shared resources should provide each user or group with the specific permissions required.

Reference: **Consequences for Files Under the APPROOT Hierarchy**

Care must be taken when defining default permissions under APPROOT directories for application administrators who will create files that need to be accessed by regular users. When application files are written under any application, they are created according to the rules of ownership and default permissions applicable to the user who signed in to the Reporting Server. Thus, application developers must share a common default group if they are to work on a shared project.

Reference: **Consequences for Files Under the EDATEMP Hierarchy**

In a server that runs with security OPSYS, special conditions apply to the edatemp subdirectory and its contents. The basic principle is that agent subdirectories under edatemp are owned at any given moment by the user that the agent is impersonating.

With an OPSYS provider, this means that once a user connects to an agent, the user is explicitly given ownership of their agent subdirectory, regardless of which user the agent impersonated during previous connections. This ownership and the inherited permissions defines how other users can access the connected user's temporary files, if at all.

On UNIX, z/OS (HFS deployment), and IBM i, the edatemp subdirectory itself is an exception because it is set up similarly to the /tmp directory. While it allows any user to add files and directories underneath, only the creator of a file can later rename or delete it. Files directly under edatemp (for example, traces) and listener subdirectories are owned by the super user (root, QSECOFR). If edatemp is pre-created or adjusted to an alternate permission, then that permission will be respected. However, be aware that edastart -cleardir will remove the edatemp directory, and therefore necessitate the repeated recreation to the desired permission.

On z/OS (PDF deployment), all temporary files created by a user are based on the rules in the Security package.

On Windows, the permissions of the parent directory of edatemp (usually EDACONF), which were applied when the server was installed, are inherited when edatemp and then the agent subdirectories are created, defining the kind of access other users have to an agent subdirectory and its contents.


Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions

An authenticated user interacts with the server using the Reporting Server browser interface or Data Migrator desktop interface. The Server Administrator controls access to the Reporting Server browser interface actions permitted by assigning general privileges to roles and registering users and groups into roles.

How to Configure General Privileges

General privileges define access to the Reporting Server browser interface and Data Migrator desktop interface control pages, such as the Adapter and DBMS Connections configuration pages, the Metadata creation pages, the Procedure editor and Procedures and Flows Run Options page, and the Server Configuration pages.

These privileges can be reviewed or customized from the Reporting Server browser interface on the *Access Control* page by right-clicking a role, user, or group and clicking *General Privileges*. General privileges can be set on the role, group, and user levels. Groups and users registered under any role inherit the general privileges of that role unless they are customized on the group or user level.

Role: Application Administrator 			
Code	Effective	Default	Privilege Description
- Adapters and Core Server <input type="checkbox"/>			
ADPTP	<input type="checkbox"/>		Configure Data Adapters for All Users
NODPT	<input checked="" type="checkbox"/>		Disable Direct Passthru
NOSYS	<input checked="" type="checkbox"/>		Disable Operating System Commands
- Metadata <input type="checkbox"/>			
METAP	<input checked="" type="checkbox"/>	✓	Launch Tools to Create/Edit/Upload Metadata
METUP	<input type="checkbox"/>	✓	Launch Tools to Upload Metadata
DATMG	<input checked="" type="checkbox"/>	✓	Data Management (Allows Users to Create DBMS Tables for the Synonym, Use Quick Copy, Upload to Relational Database and Re-load Data)
- Procedures <input type="checkbox"/>			
PRSAV	<input checked="" type="checkbox"/>	✓	Launch Tools to Edit Procedures, Upload Procedures
PRSTR	<input type="checkbox"/>		Stress Tool
PRDFR	<input checked="" type="checkbox"/>	✓	Schedule/E-Mail/Submit
PRRPT	<input checked="" type="checkbox"/>	✓	View Output, DM Log and Statistics, Impact Analysis, Scheduler and Flow Reports
PROUT	<input checked="" type="checkbox"/>	✓	View Log and Output of Submitted Requests
- Workspace <input type="checkbox"/>			
WSCFG	<input type="checkbox"/>		Server Administrative functions(Access Control, Server Configuration and Migration, Scalability, Change Application Path)
GRANT	<input type="checkbox"/>		Grant Permissions
MONIT	<input checked="" type="checkbox"/>	✓	Monitor Agents, Sessions, Connections and Services
KILAL	<input type="checkbox"/>		Kill and Stop Agents, Sessions, Connections and Services

Users other than the Server Administrator can see their own general privileges by choosing *Show My General Privileges* on the My Console menu. Server Administrators do not have this option because they have full privileges to all pages of the Reporting Server browser interface and Data Migrator desktop interface.

WebFOCUS procedures can use the CHECKPRIVS() function that, when given a privilege code (for example, NODPT), returns the value Y (yes) or N (no) depending on whether the connected user has that privilege. For more information, see the *TIBCO WebFOCUS® Using Functions* manual.

Reference: General Privileges

The following table describes the general privileges.

Privilege	Description
Adapters and Core Server	

Privilege	Description
ADPTP	You can configure adapters and add DBMS connection attributes.
NODPT	You can disable Direct Passthru.
NOSYS	You can disable the ability to execute certain operating system commands.
Metadata	
METAP	You can launch tools to create and edit metadata.
METUP	You can create synonyms during upload, quick copy, and custom copy.
DATMG	Data Management (allows users to create DBMS tables for the synonym, run Quick Copy, upload to Relational Database, and re-load data).
Procedures	
PRSAV	You can launch tools to edit procedures and upload procedures.
PRSTR	You can use the stress tool.
PRDFR	You can Schedule/Email/Submit.
PRRPT	You can view output, Data Migrator log and statistics, impact analysis, scheduler and flow reports.
PROUT	You can view log and output of submitted requests.
Workspace	
WSCFG	You can perform server administrative functions, such as access control, server configuration and migration, scalability control, and application path control.
GRANT	You can grant file privileges that you have to other users.
MONIT	You can monitor agents, sessions, connections, and services for all users.

Privilege	Description
KILAL	You can kill and stop agents, sessions, connections, and services for all users.
SRVLG	You can view server logs and traces, and you can create Savediag.
STPSV	You can stop and restart the server.
RARGP	You can use Resource Analyzer and Resource Governor.
My Console	
CHGPW	You can change your password.
MONUS	You can monitor Data Service Agents that match your user ID.
MONGR	You can monitor Data Service Agents that match your group ID.
KILT3	You can kill Data Service Agents that match your user ID.
KILGR	You can kill Data Service Agents that match your group ID.
APATH	You can change your application path (no applock).
UPROF	You can edit your user profile.
DBMSC	You can manage your own DBMS connections.
UPROF	You can edit your own user profile.
APROF	You can list application profiles.
UINFO	You can disable display of My Console, detailed error messages, sign-in information, list of privileges for the user, server version, console log, help, and the Privileges section of the Properties page for a file or directory.

Procedure: How to Control Access to the Reporting Server Browser Interface and Data Migrator Desktop Interface

Authenticated users are permitted to perform certain control operations on the Reporting Server browser interface, depending on their definitions in the `admin.cfg` file as Server Administrator, Application Administrator, Server Operator, Basic User, or other custom role defined on the server. Within a given role, additional administration privileges may be applied. As a result, available Reporting Server browser interface features may vary from site to site or between server configurations at a site.

The Reporting Server browser interface or Data Migrator desktop interface functions available to the user are determined by their role. If a user is not assigned a role, does not belong to a group, or the group is not registered, their access to Reporting Server browser interface functions is based on the `default_admin_role` setting.

Access to the Reporting Server browser interface is also affected by the security provider in effect:

Security OPSYS. Access to the Reporting Server browser interface is protected by user authentication at the operating system level.

Security PTH. Access to the Reporting Server browser interface is protected by user authentication against the `admin.cfg` file.

Security DBMS. Access to the Reporting Server browser interface is protected by user authentication against the selected DBMS.

Security LDAP. Access to the Reporting Server browser interface is protected by user authentication against the LDAP (or AD) server.

Security OFF. With security OFF, anyone can access the Reporting Server browser interface, with full, unrestricted use of its features.

With security on, if a server administrator has established an Anonymous ID, it can be used to access the Reporting Server browser interface without an explicitly entered user ID. For details, see [Setting an Anonymous User ID](#) on page 59.

Access to the Reporting Server browser interface and Data Migrator desktop interface can be restricted for individual users or for a group of users.

1. Access the Reporting Server browser interface with a server administrator user ID.
2. Go to the *Access Control* page.
3. Expand a *Roles* folder.
4. Right-click a user or group and select *Properties*.

The General Properties page opens.

5. Select a role from the *Inherit Privileges from* drop-down list.

The screenshot shows a dialog box titled "Activate Providers" with a tab labeled "Group ID". The "Security Provider" field contains "PTH". The "Group" field contains "grp01_app". The "Description" field is empty. The "Inherit Privileges from" dropdown menu is open, displaying a list of roles: "Application Administrator" (highlighted), "Server Administrator", "Server Operator", "Application Administrator", "Basic User", and "None". At the bottom left, there are "Update" and "Cancel" buttons.

The options are:

- Server Administrator
- Application Administrator
- Server Operator
- Basic User
- None. (No access allowed.)

Select the value you want to apply. If only users defined in the admin.cfg file are to be allowed access, select *None*.

Note: If the Server Administrator has created custom roles, they will be available in the drop-down menu.

6. Click *Update*.

Note: If a server administrator has established an Anonymous ID, it can be used to access the Reporting Server browser interface without an explicitly entered ID. For more details, see [Setting an Anonymous User ID](#) on page 59.

Procedure: How to Minimize User Access to the TIBCO WebFOCUS Reporting Server Through the Reporting Server Browser Interface

Users need some access to the Reporting Server browser interface in order to execute procedures and upload files. However, you can minimize their access to the server when they are signed in to the Reporting Server browser interface by implementing the following steps:

1. Remove all general privileges. Set all *disable privileges*: NODPT, NOSYS, UINFO.

2. Set all file privileges for * to Execute and List.
3. Adjust privileges to applications so they only allow necessary actions.

When a user tries to perform an action that you have not permitted, a message displays advising the user to contact the WebFOCUS Administrator.

Procedure: How to Set the Default Administration Role

A default administration role provides access to Reporting Server browser interface or Data Migrator desktop interface functions and to the server for users who are not assigned a role, do not belong to a group, or whose group is not registered. It is based on a selected role.

1. Access the Reporting Server browser interface with a server administrator user ID.
2. Go to the *Access Control* page.
3. Click *Settings*, then click *Access Control*.

The Access Control Settings page opens.

4. Select a role from the *default_admin_role* drop-down menu in the General section.
5. Click *Apply and Restart Server*.

Transferring File Permissions With the GRANT Privilege


The server administrator normally provides access control permissions for application folders to groups and users. If the server administrator issues the GRANT privilege to another security subject (role, group, or user), that security subject can then grant its own file permissions to any other security subject,

For example, if user A has Read but not Write permission on folder X, the user can transfer Read permission to user B. The user cannot grant Write permissions to anyone.

Note that if user A loses the permission later on, user B will retain his transferred permission.


Example: Granting Permissions to Another User

In the following configuration, user pgmtst2 is an Application Administrator with the GRANT privilege:

User: PTH\pgmtst2 

Code	Effective	Default	Privilege Description
- Adapters and Core Server	<input type="checkbox"/>		
ADPTP	<input type="checkbox"/>		Configure Data Adapters for All Users
NODPT	<input checked="" type="checkbox"/>	✓	Disable Direct Passthru
NOSYS	<input checked="" type="checkbox"/>	✓	Disable Operating System Commands
- Metadata	<input type="checkbox"/>		
METAP	<input checked="" type="checkbox"/>	✓	Launch Tools to Create/Edit/Upload Metadata
METUP	<input type="checkbox"/>		Launch Tools to Upload Metadata
DATMG	<input checked="" type="checkbox"/>	✓	Data Management (Allows Users to Create DBMS Tables for the Synonym, Use Quick Copy, Upload to Relational Database and Re-load Data)
- Procedures	<input type="checkbox"/>		
PRSAV	<input checked="" type="checkbox"/>	✓	Launch Tools to Edit Procedures, Upload Procedures
PRSTR	<input type="checkbox"/>		Stress Tool
PRDFR	<input checked="" type="checkbox"/>	✓	Schedule/E-Mail/Submit
PRRPT	<input checked="" type="checkbox"/>	✓	View Output, DM Log and Statistics, Impact Analysis, Scheduler and Flow Reports
PROUT	<input checked="" type="checkbox"/>	✓	View Log and Output of Submitted Requests
- Workspace	<input type="checkbox"/>		
WSCFG	<input type="checkbox"/>		Server Administrative functions(Access Control, Server Configuration and Migration, Scalability, Change Application Path)
GRANT	<input checked="" type="checkbox"/>		Grant Permissions
MONIT	<input checked="" type="checkbox"/>	✓	Monitor Agents, Sessions, Connections and Services
KILAL	<input type="checkbox"/>		Kill and Stop Agents, Sessions, Connections and Services

User pgmtst2 has Read, Execute, and List privileges on application app06, while user pgmtst3 has no privileges on this application:

File/Directory: C:\ibi\apps\app06 Customized 

Subject	Read	Write	Execute	List	Inherit Subject	Inherit Directory	Description
+ Server Administrator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Administrator	*	
- Server Operator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Operator	C:\ibi\apps	
User - PTH\pgmtst3	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Customized		
User - PTH\svop1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Operator	C:\ibi\apps	
- Application Administrator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Application Administrator	C:\ibi\apps	
User - PTH\pgmtst2	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Application Administrator	C:\ibi\apps	
User - PTH\svvaa1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Application Administrator	C:\ibi\apps	
Group - PTH\grp01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Application Administrator	C:\ibi\apps	
Basic User	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Basic User	C:\ibi\apps	
None	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - None	*	

Since user pgmtst2 has the GRANT privilege for this application, user pgmtst2 can edit the privileges for user pgmtst3.

User: PTH\pgmtst3 Customized [Reset](#)

Directory/File	Read	Write	Execute	List	Status	Inherit Subject	Inherit Directory	Remark
-*	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Inherited	Role - Server Operator	*	
- C:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Inherited	Role - Server Operator	*	
- ibi	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Inherited	Role - Server Operator	*	
- apps	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Inherited	Role - Server Operator	C:\ibi\apps	APPROOT
app06	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Customized			

Configuring Groups

The Server Administrator registers a group or user within a role that has a set of privileges. Users and groups can be registered for special privileges, both general privileges and file privileges. It is a good practice to assign users to groups and to control groups rather than individual users. The installation has a choice of grouping users in LDAP or using the operating system security mechanism. This grouping is external to the WebFOCUS or server software and must be accomplished by the administrators for those external products. The groups are then registered on the server, and a role is chosen for them. By default, the privileges are inherited from the role in which the group is registered.

At sign-in time, the user group is determined and the proper privileges are assigned. If the user does not belong to any registered group or role, the privileges are taken from the default role. At installation time, the default role is set to Basic User, but the server administrator can change this designation. The default role can be reset on the Access Control General Privileges pages using the default_admin_role keyword.

A group can be assigned a profile stored as *groupname.prf*. Group profiles support the same syntax as role profiles, which are described in [Configuring Roles](#) on page 85.

The right-click menu for a user or group enables the Server Administrator to view or edit the Application Path, General Privileges, and Directory/File Privileges. It also has selections that show, for a user, the list of group memberships, and, for a group, the list of members.

Note: Profile settings affect how groups connect to the server.

Note that for LDAP security, if the LDAP database contains the email address, it is automatically populated. This requires the ldap_user_email attribute to be set in the LDAP provider configuration.

If a user belongs to multiple groups, the server either uses the primary group privileges, allows the user to select the effective group at sign-in time, or merges the privileges from all groups. How the server handles privileges for users in multiple groups is determined by the `group_profile` setting in the `edaserve.cfg` file. You can access this setting from the Access Control page by right-clicking the top-level of the Access Control tree and selecting *Access Control*. The default value for `group_profile` setting is to use the primary group. If the security provider is OPSYS, the operating system determines the primary group. For LDAP, the primary group is the first group alphabetically for the user.

GRPLIST is a function that returns a group name or a list of group names (separated by colons) for the connected user. This function is supported for LDAP security with all types of connections. If the group list is empty or there is an error in the function parameters, the function returns blanks.

There are also three &vars (`&FOCSECUSER`, `&FOCSECGROUP`, and `&FOCSECGROUPS`) that store the connected user and its primary group or the list of all groups (if the server is configured with `profile_setting=all`). These variables are populated by the server and cannot be changed by the application.

Customizing Group Privileges

Although the recommended policy is to control groups, there may be a need to create different sets of privileges for some individual users. If a user cannot be assigned to a group or to a custom role, the user can be registered with any role, and the Server Administrator can customize the privileges as needed. The customization can include general and file privileges, application path, and DBMS connections. Such users should be kept to a minimum.

Configuring Roles

The server is shipped with five predefined roles that allow basic operational control. Roles are listed on the *Access Control* page:

The roles come with a fixed set of general privileges as follows:

- Server Administrator.** Groups and users assigned to this role have full control of the workspace, adapters, synonyms, and procedures.
- Application Administrator.** Groups and users assigned to this role have privileges to create synonyms and procedures and to run them.
- Server Operator.** Groups and users assigned this role have privileges to start and stop the server and monitor and kill agents.
- Basic User.** Groups and users assigned to this role can execute procedures.

- None.** Users assigned to this role have no privileges on the Reporting Server browser interface and Data Migrator desktop interface.

The default file and directory privileges allow full access.

The Server Administrator can customize the default privileges and can create custom roles and assign them general and file privileges. The only roles that cannot be changed are the Server Administrator role, which always has full privileges, and the NONE role, which has no privileges. If the Server Administrator customizes a role, it displays with the text (Customized) next to its name on the Access Control tree.

Roles can be assigned profiles that are stored in the location identified by the \$EDAPRFU variable. The profile for each role is stored as the file named rolename.prf. The following items can be controlled by a profile:

- DBMS connections that are configured on the Adapters page.
- Application path that controls the search for unqualified file names (for example, TABLE FILE CAR). The application path is configured from the Access Control page.
- Other WebFOCUS commands.

Predefined roles are assigned predefined names:

Name	Description
SRV	Server Administrator
APP	Application Administrator
OPR	Server Operator
USR	Basic User

Note: When a user connects to the server, the role profile is executed after the server profile.

Procedure: How to Create a Custom Role

1. From the *Access Control* page, either click the *Register* button at the top of the page and select *Register Role*, or right-click the Roles folder, and click *Register Role* from the context menu.

The *Register Custom Role* page opens.

2. Enter a name for the role, and click *Apply*.

A dialog box opens informing you that a new role will be registered.

3. Click *OK*.

The *Activate Providers* page opens.

4. You can assign privileges by right-clicking the role and selecting *Directory/File Privileges* or *General Privileges* from the context menu and checking the privileges you want to assign to this role.
5. Click *Save* to save the privileges.

Note:

- Once the custom role has been created, the Server Administrator can register users and groups to the role.
- The Server Administrator can create a role profile that will be executed when a user registered to this role connects to the server.
- The Server Administrator can set the custom role as `default_admin_role` on the *Access Control/General Settings* page.

Procedure: How to Set Customized Privileges for Roles

A server administrator can set selected privileges for basic users, application administrators, and server operators. Configurable options are tailored to each user group.

1. On the Reporting Server browser interface menu bar, select *Access Control*.
2. Right-click a role in the navigation pane and select *General Privileges*.

The *General Privilege* page opens.

3. Select the check boxes for the functions you want the role to be able to perform.
Configurable options vary by user group.
4. Click *Save* to confirm these settings.
5. Click *Reset to Default Privileges* if you wish to revert to the standard privileges assigned to that user group.

Registering Users and Groups in a Role

A server administrator can assign specific users the following roles: Server Administrator, Application Administrator, Server Operator, Basic User, or None. A server administrator can also create additional roles and select or deselect certain privileges for these users or groups.

The user and group roles are stored in a release independent file called `admin.cfg`, which is located, by default, in `.../ibi/profiles`. This location is defined with the server configuration parameter `EDAPRFU`, which is stored in the `edaserve.cfg` file. The `admin.cfg` file should not be shared between servers that run with different security providers.

When registering users and groups under an LDAP or OPSYS security provider, you can retrieve a list of groups and users from the operating system. The list can be filtered to minimize the number of candidates retrieved. The User and Group Description is retrieved when available. This feature is available on Windows, UNIX, iSeries, and zOS/USS.

Procedure: How to Register Users

This procedure assumes that the user is signed in to the Reporting Server browser interface as a server administrator.

1. From the *Access Control* page, right-click a role and select *Register User* from the context menu or click *Register*, then *Register User* from the ribbon.

The *User Registration* page opens, as shown in the following image.

By default, the *Import From Provider* button is selected.

This enables you to select a security provider and retrieve a list of users whose attributes (user ID, Description, and Email) match a pattern, and select one or more from the list to register. If you do not enter a user ID or a pattern, all users are retrieved from the security provider directory.

2. To import users from a security provider directory, select the security provider from the drop-down list, and enter values or patterns in one or more of the input fields (User ID, Description, and EMail), and click *Next*.

A list of retrieved users opens, as shown in the following image.

<input type="checkbox"/>	User	Description	E-Mail	Registered Role
<input type="checkbox"/>	pmsaep			
<input type="checkbox"/>	pmsesd			
<input type="checkbox"/>	pmsjas			
<input type="checkbox"/>	pmsmjb			
<input type="checkbox"/>	pmssae			

- a. Click the check box next to each user you want to register.
- b. If you want to change the role from which the users will inherit privileges, select a role from the *Inherit Privileges from* drop-down list.

- c. If you want to override an existing registration for the selected users, check *override the existing registration*.
 - d. Click *Register* to register the users or *Back* to go back to the previous page.
3. For single-user registration, click *Manual*.

The Register User page displays with the fields needed for registering a user for providers for which the server does not have support for retrieving user and group lists from the security provider. For example, single registration is used for a DBMS provider or an OPSYS provider on z/OS PDS deployment, as shown in the following image.

If you register a user with a provider that has its own user and group directory, and the user or group you register is not in that directory, the user or group will not be able to be authenticated.

- 4. Select the security provider for the registration from the *Security Provider* drop-down list.
- 5. Enter the following properties.
 - Enter a user ID in the *User* field.
 - If you use domain names with the user IDs (Windows only), enter the domain name in the *Domain* field or select one from the drop-down list (OPSYS only).
 - Optionally, enter a description in the *Description* field.
 - Optionally, enter an email address for the user.
 - Select an administrator level from the *Inherit Privileges from* drop-down list. The default options are: *Server Administrator*, *Application Administrator*, *Server Operator*, *Basic*, or *None*.

- Optionally, in the *Password* and *Confirm Password* fields, enter the password for the user ID. A password is only required for scheduled runs.
- Click *Add and Register*.

Procedure: How to Register Groups

This procedure assumes that the user is signed in to the Reporting Server browser interface as a server administrator.

Note that for the PTH security provider, the group first has to be created using the PTH Users and Groups Management page.

For the OPSYS, LDAP, Custom, or PTH security providers, a server administrator can assign a role to groups of users by following these steps:

1. From the *Access Control* page, right-click a role and select *Register Group* from the context menu, or click *Register*, then *Register Group* from the ribbon.

The *Group Registration* page opens, as shown in the following image.

By default, the *Import From Provider* button is selected.

This enables you to select a security provider and retrieve a list of groups whose attributes (user ID and Description) match a pattern, and select one or more from the list to register. If you do not enter a group ID or a pattern, all groups are retrieved from the security provider directory.

2. To import users from a security provider directory, select the security provider from the drop-down list, and enter values or patterns in one or more of the input fields (User ID and Description), and click *Next*.

A list of retrieved groups opens, as shown in the following image.

The screenshot shows the 'Group Registration' page. At the top, there are two tabs: 'Activate Providers' and 'Group Registration'. Below the tabs, there are two fields: 'Security Provider' set to 'LDAP - Idap01' and 'Inherit Privileges from' set to 'Application Administrator'. There is a checkbox for 'override the existing registration' which is currently unchecked. Below this are two buttons: 'Back' and 'Register'. At the bottom, there is a table with the following columns: 'Group', 'Description', and 'Registered Role'. The table contains four rows of data, each with a checkbox in the first column.

<input type="checkbox"/>	Group	Description	Registered Role
<input type="checkbox"/>	ats	Advanced Technical Support	
<input type="checkbox"/>	core_admin	Core Engineering Administrators	
<input type="checkbox"/>	dev	Server Development	
<input type="checkbox"/>	Directory Administrators		

- a. Click the check box next to each group you want to register.
 - b. If you want to change the role from which the groups will inherit privileges, select a role from the *Inherit Privileges from* drop-down list.
 - c. If you want to override an existing registration for the selected groups, check *override the existing registration*.
 - d. Click *Register* to register the groups or *Back* to go back to the previous page.
3. For single group registration, click *Manual*.

The *Group Registration* page displays with the fields needed for registering a group for providers for which the server does not have support for retrieving user and group lists from the security provider. For example, single registration is used for a DBMS provider or an OPSYS provider on z/OS PDS deployment, as shown in the following image.

The screenshot shows the 'Group Registration' page with the 'Manual' radio button selected. The 'Security Provider' is set to 'LDAP - Idap01'. There are three input fields: 'Group', 'Description', and 'Inherit Privileges from' (set to 'Application Administrator'). A 'Register' button is at the bottom.

If you register a group with a provider that has its own user and group directory, and the user or group you register is not in that directory, the user or group will not be able to be authenticated.


4. Select the security provider for the registration from the *Security Provider* drop-down list.
5. Enter the following properties.
 - Enter a group ID in the *User* field.
 - If you use domain names with the user IDs (Windows only), enter the domain name in the *Domain* field or select one from the drop-down menu (OPSYS only).
 - Optionally, enter a description in the *Description* field.
 - Select an administrator level from the *Inherit Privileges from* drop-down list. The default options are: *Server Administrator*, *Application Administrator*, *Server Operator*, *Basic*, or *None*.
 - Click *Register*.

***Procedure:* How to Manage Users and Groups With PTH Security**

This procedure assumes that the user is signed in to the Reporting Server browser interface as a server administrator.

1. Go to the *Access Control* page.
2. Expand *Security Providers*.
3. Right-click the PTH provider and click *Manage Users/Groups*.

The *PTH Users and Groups Management* page opens, as shown in the following image.



4. You can drag users (listed on the left) to groups (listed on the right) or drag groups to users.

Note: You must click Save in order to make these changes permanent.

5. To add a user or group, click New on the appropriate pane.

For a new user, the New PTH User dialog box opens, as shown in the following image.

- a. Enter a user ID.
- b. Optionally, enter a description in the *Description* field.
- c. In the *Password* and *Confirm Password* fields, enter a password for the user ID.
- d. Enter an email address for the user.
- e. By default, the password never expires. If you want to set an expiration date, uncheck *Password never expires*.
- f. If you want to temporarily disable the account, check *Account disabled*.
- g. Click *OK*.

The user is added to the list.

- h. Click *Save*.

Note: If you do not click *Save*, the user will not be saved on the list.

For a new group, the New PTH Group dialog box opens, as shown in the following image.

- a. Enter a group ID.
- b. Optionally, enter a group description.
- c. Click *Next*.

The group is added to the list.

- d. Click *OK*.

Note: If you do not click *Save*, the group will not be saved on the list.

6. To delete a user or group, select the ID and click *Delete*.

A confirmation dialog box opens.

- a. Click *OK*.
- b. Click *Save*.

Note: If you do not click *Save*, the user or group will not be deleted.

7. To view or change the properties of a user or group, select an ID and click *Properties*.

For more information, see [How to View or Change Existing User or Group Registrations](#) on page 96.

Procedure: How to View or Change Existing User or Group Registrations

1. From the *Access Control* page, expand a Role from the *Roles* folder.
2. Right-click the user or group whose properties you want to view or change, and select *Properties*.

The *Properties* page opens, as shown in the following image.

? Security Provider	PTH
? User	pgmtst2
? Description	
? E-Mail	
? Inherit Privileges from	Application Administrator

Update Cancel

3. You can change the user role by selecting a different role from the *Inherit privileges from* drop-down list. For more information, see [Configuring General TIBCO WebFOCUS Reporting Server and Reporting Server Browser Interface Actions](#) on page 76.

Note: If only one user is defined as a Server Administrator, you cannot change that user role.

4. Click *Update*.

The user or group is moved to the new role.

Procedure: How to Unregister a User or Group

1. From the Access Control page, expand a Role from the *Roles* folder.
2. Right-click the user or group and select *Unregister*.

You are prompted to confirm that you want to unregister the user or group.

3. Click *OK*.

The user or group is unregistered.

Configuring Encryption

The server supports encryption of passwords in configuration files, as well as SSL encryption for the TCP/HTTP Listener and encryption of data passed between a hub server and a remote server or cluster server.

Encrypting Passwords Stored in Configuration Files

You can select an encryption algorithm for passwords stored in configuration files.

To encrypt passwords stored in configuration files, do the following:

1. From the *Access Control* page, click *Settings*, then *Encryption*.

The Encryption and Decryption page opens.

2. To define the cipher used to encrypt passwords in configuration files, click the drop-down list for *cfgfile_cipher* and select:

- DES* for Data Encryption Standard.
- 3DES* for Triple Data Encryption Standard.
- AES 128 bits* for Advanced Encryption Standard (key size 128 bits). This is the default encryption algorithm.
- AES 256 bits* for Advanced Encryption Standard (key size 256 bits).
- user defined program* to specify your own encryption and decryption programs.

3. Click *Apply and Restart Server*.

User-Defined Password Encryption and Decryption

You can configure the Reporting Server to employ user-defined password encryption and decryption programs.

User-defined encryption and decryption are available from the *Encryption* option on the Reporting Server browser interface Access Control resources tree or ribbon. Set the parameter *cfgfile_cipher* to *user defined program*, which enables you to specify your password encryption and decryption programs by entering them in the *cfgfile_cipher_encrypt* and *cfgfile_cipher_decrypt* text boxes.

Note: If you choose to use password encryption outside the server, only the path to the decryption program needs to be specified.

In cases where a password is encrypted outside of the server, the encrypted password string (or label) must be used, instead of the password in server and adapter configurations. When user-defined program is chosen, all passwords for registered users in the *admin.cfg* file, all adapter connections in the *edasprof.prf* file (or *user*, *group* and *role* profiles), passphrases in the *odin.cfg* file, and pooled user passwords need to be encrypted using the same encryption program. They will be decrypted using the same decryption program. Remote Server connections in the *odin.cfg* file do not support the user-defined encryption option. If user-defined encryption is chosen, the Cluster Manager feature should not be enabled.

Set Up User-Defined Password Encryption and Decryption

1. On the Access Control page, click *Settings*, then *Encryption*.
2. Select *user defined program* from the *cfgfile_cipher* drop-down menu.

Two additional parameters are displayed.

3. Enter the full paths to the encryption and decryption programs in the *cfgfile_cipher_encrypt* and *cfgfile_cipher_decrypt* fields respectively, as shown in the following image.

The screenshot shows a configuration window with the following fields and values:

- cfgfile_cipher**: user defined program
- cfgfile_cipher_encrypt**: (empty text field)
- cfgfile_cipher_decrypt**: (empty text field)
- io_encryption**: n
- encrypt_admincfg**: n

At the bottom of the window are two buttons: **Cancel** and **Save**.

4. Click **Save**.

Encrypting Data at Rest

When the server *io_encryption* setting is ON, binary, alpha, and delimited HOLD files in *edatemp* and *foccache* are encrypted. Data agent trace files (*tsxxx.trc*) and agent output files (*tsxxx.tro*) are also encrypted as a result of this setting. After encryption, the extensions for these files are *.trce* and *.troe*, respectively. This prevents the user from opening these files in editors outside of the Reporting Server browser interface or Data Migrator desktop interface. Trace files and agent output files can be viewed from the Reporting Server browser interface *Workspace* page under *Traces And Logs*.

To turn on the *io_encryption* setting, do the following:

1. On the Access Control page, click *Settings*, then *Encryption*.
2. Select *y* from the *io_encryption* drop-down list.

The default encryption algorithm is *AES 128 bits*. You can change the type of encryption using the `cfgfile_cipher` setting available on the same page.

3. Click **Save**.

This setting will be added in the `edaserve.cfg` configuration file.

Configuring Secure Socket Layer (SSL) Encryption for the TCP/HTTP Listener

You can enable SSL for the TCP/HTTP Listener to encrypt all traffic between the server and any client application, such as the WebFOCUS Client, a remote server, or a cluster server.

To enable SSL, do the following:

1. On the *Workspace* page, open the *Special Services and Listeners* folder, right-click *TCP/HTTP*, and select *Properties of HTTP*.

The Listener Configuration page opens.

The screenshot shows the 'Security' section of the Listener Configuration page. The 'Enable HTTPS' dropdown menu is open, showing options: 'No', 'Yes (OpenSSL)', and 'Yes (Microsoft)'. The 'No' option is currently selected. Other settings in the Security section include: 'RESTRICT_TO_IP' (empty text field), 'LOGIN_FAILURE' (0), 'LOGIN_HIDE_PROVIDERS' (0), 'CSRFS_TOKEN' (1), 'SANITIZE_PARAMETERS' (1), 'VERIFY_REFERER' (1), and 'X_FRAME_OPTIONS' (SAMEORIGIN). The 'Sessions Control' section is set to 'No'. The 'Aliases' and 'Miscellaneous Settings' sections are collapsed. At the bottom, there are 'Save and Restart Server' and 'Cancel' buttons.

2. In the *Security* section, select *Yes (OpenSSL)* or *Yes (Microsoft)* from the *Enable HTTPS* drop-down list.

Additional fields needed for SSL configuration appear.

Note:

- ❑ Open SSL libraries v1.1.1 and v1.0.2 are supported.

The current run-time version depends on the following configuration keyword in the `edaserve.cfg` file.

```
ssl_lib_vendor = {sslv11|sslv10}
```

The default value for `ssl_lib_vendor` is `sslv11`.

You should also set the environment variable `IBI_LOAD_SSL_FROM` or `IBI_LIBPATH` to point to the appropriate openSSL library.

For example:

```
ssl_lib_vendor = sslv11
IBI_LIBPATH=/usr/local/ssl64/1.1.1/lib
```

or:

```
ssl_lib_vendor = sslv10
IBI_LIBPATH=/usr/local/ssl64/1.0.2/lib
```

- ❑ OpenSSL libraries `libeay32.dll` and `ssleay32.dll` must be in the path to enable SSL.

3. Enter the following values:

`SSL_CERTIFICATE`

Contains the certificate chain in order, starting with the certificate for the listener and ending with the root CA certificate. Each of these entries must be in PEM format.

Note that the administrator at the installation site must acquire valid security certificates (self-signed or commercial).

`SSL_PRIVATE_KEY`

Defines the file that contains the private key of the listener. It must correspond to the public key embedded within the certificate and must be in PEM format.

`SSL_PASSPHRASE_E`

If the file defined in `SSL_PRIVATE_KEY` is encrypted, a passphrase must be provided here to decrypt the private key.

`SSL_CA_CERTIFICATE`

Defines the name of a file containing a trusted CA certificate in PEM format. It is used to verify the client certificate. If the client fails to send a certificate or verification fails, connections are rejected. More than one CA certificate may be present in the file.

`SSL_LABEL`

Specifies a descriptive comment used to identify the key and certificate in the certificate database. (IBM GSK Only).

For z/OS only, is a certificate LABEL in the RACF database. Must be connected to the RACF Keyring pointed by the parameter SSL_CERTIFICATE. Server system ID (default ISERVER) must be the certificate owner in the case when a personal certificate is used. When a SITE certificate is used as a server certificate, it has to be connected to a Keyring with the parameter USAGE(PERSONAL). Otherwise, the private key is unaccessible.

SSL_FRIENDLY_NAME

Specifies a name used to identify the certificate in the PKCS#12 file, if the file contains more than one certificate.

SSL_AUTH_CLIENT

Defines whether client authentication should be enabled. Select one of the following values from the drop-down list.

- 0.** The client is not asked for a certificate. There is no client authentication.
- 1.** The server requests the client to send a certificate. All received certificates are validated. If the client fails to send a certificate or verification fails, connections will be rejected.

4. Click *Save and Restart Server*.

Support for OpenSSL Protocols and Ciphers, Microsoft Algorithm Identifiers, and DH and ECDH Key Exchange

The following parameters are supported in the edaserve.cfg configuration file to support OpenSSL protocols and ciphers, Microsoft algorithm identifiers, and Diffie-Hellman (DH) and Elliptic Curve Diffie-Hellman (ECDH) key exchange.

ssl_protocol

Defines the acceptable transport layer protocols. The default is:

```
ssl_protocol = TLS1_2
```

Note: There are no spaces after the commas.

To disable a particular protocol, you remove it from the list of supported protocols provided in the ssl_protocol keyword.

ssl_ciphers = *cipher strings* (OpenSSL)

Configures the list of supported cipher strings for OpenSSL. For example, the following is the default.

```
ALL:!ADH:!LOW:!EXP:@STRENGTH
```

`ssl_ciphers = algorithm identifiers (Microsoft)`

When using Microsoft cryptography, you can specify algorithm identifiers, effectively limiting ciphers utilized in HTTPS connections.

The following are supported algorithm identifiers (ALG_ID):

`CALG_SHA_384,CALG_SHA_256,CALG_SHA,CALG_AES_256,CALG_AES_128,
CALG_DH_EPHEM,CALG_RSA_SIGN,CALG_RSA_KEYX`

The identifiers should be comma-separated with no spaces. For example:

`ssl_ciphers = CALG_SHA_384,CALG_SHA_256,CALG_AES_128`

Invalid algorithm identifiers are ignored.

`ssl_dhparams_file`

Specifies the full path to the Diffie-Hellman (DH) parameters file. To create a DH key exchange with key size 1024 bits, issue the following command.

`'openssl dhparam -outform PEM -out dhParam.pem 1024'`

`ssl_ecdhcurve`

¥

Specifies a curve name for a specific Elliptic Curve Diffie-Hellman key exchange. To list all available curves, issue the following command.

`'openssl ecparam -list_curves'`

Configuring Data Encryption for a Remote TIBCO WebFOCUS Reporting Server

You can enable encryption of data passed between the server and a remote server or cluster server.

1. You can access the Remote Server Configuration page from the *Get Data* page. Remote Servers is an item on the *New Datasource* list.
2. Right-click *Remote Servers* on the list and click *Add Remote Server Connection*.

The Remote Server Configuration page opens.

The screenshot shows a configuration window with two tabs: 'Basic' and 'Advanced'.
Basic Tab:
 ? NODE: node00 (type in a new one or select one from the list, character ';' is not allowed)
 ? HOST: (empty field) (required)
 ? PORT: 8100 (required)
 ? HTTP_PORT: 8101
 ? SECURITY: Trusted
 ? profile: edasprof (type in a new one or select one from the list)
 ? Windows User: (empty field) for server on Windows only
 ? Windows Password: (empty field) for server on Windows only
 ? Windows Service: (empty field) for server on Windows only
 ? DESCRIPTION: (empty field) character ';' is not allowed
Advanced Tab:
 ? SERVICE_NAME: (empty field)
 ? COMPRESSION: 0
 ? ENCRYPTION: 0 (dropdown menu is open showing options: 0, DES, ADVANCED, IBCRYPT)
 ? CONNECT_LIMIT: (empty field)
 ? MAXWAIT: (empty field)
 Buttons: Save, Cancel

3. Click the *ENCRYPTION* drop-down list and select:

- 0 for no encryption.
- DES* for 56-bit fixed-key Data Encryption Standard in Electronic Code Book (ECB) mode. The same key is used in all connections with no key exchange between client and server.
- ADVANCED* to select an encryption cipher (3DES, AES128, AES192 or AES256), encryption mode (ECB or CBC), and RSA key length (512 or 1024 bits). In advanced mode, the client randomly generates a new RSA key pair (public and private keys of the specified length) and sends the public key to the server. Upon receipt of the public key, the server generates a random secret key. The length of the secret key depends on the chosen cipher strength. The secret key is encrypted with the public RSA key and sent back to the client, which decrypts it with its private RSA key. After the exchange, the client and the server both share the same secret key, and use it to encrypt and decrypt all communications between them.

The following encryption ciphers are available:

- 3DES for triple Data Encryption Standard.
- AES128 for Advanced Encryption Standard (key size 128 bits).
- AES192 for Advanced Encryption Standard (key size 192 bits).
- AES 256 bits for Advanced Encryption Standard (key size 256 bits).

The following encryption modes are available:

- ECB for Electronic Code Book mode. This is the default mode.
- CBC for Cipher Block Chaining mode.

The following RSA key lengths are available:

- 512 bits.
- 1024 bits.
- IBCRYPT for a user-defined algorithm. The key is 512-bit RSA-encrypted.

4. Click Save.

To enable remote WebFOCUS Reporting Servers running on the z/OS or AS/400 platform to access cryptographic services through the SESSION ONLY TOKEN, you must give the remote server ID or a remote server group RACF permissions.

To give access to CSFSERV class facilities to an arbitrary GRP group, add permissions as shown in the following example:

```
PERMIT CSF1GKP CL(CSFSERV) ID(GRP)
PERMIT CSF1TRD CL(CSFSERV) ID(GRP)
PERMIT CSF1PKV CL(CSFSERV) ID(GRP)
PERMIT CSF1DVK CL(CSFSERV) ID(GRP)
PERMIT CSF1GAV CL(CSFSERV) ID(GRP)
PERMIT CSF1PKS CL(CSFSERV) ID(GRP)
PERMIT CSF1SKD CL(CSFSERV) ID(GRP)
PERMIT CSF1SKE CL(CSFSERV) ID(GRP)
PERMIT CSF1TRC CL(CSFSERV) ID(GRP)
PERMIT CSF1TRD CL(CSFSERV) ID(GRP)
```

To give access to the CRYPTOZ class to an arbitrary GRP group, add permissions as shown in the following example:

```
RDEF CRYPTOZ CLEARKEY.SYSTOK-SESSION-ONLY UACC(NONE)
PERMIT CLEARKEY.SYSTOK-SESSION-ONLY CLASS(CRYPTOZ) ID * ACC(UPDATE)
RDEFINE CRYPTOZ SO.GRP.* UACC(NONE)
RDEFINE CRYPTOZ USER.GRP.* UACC(NONE)
PERMIT USER.GRP.* CLASS(CRYPTOZ) ID * ACCESS(UPDATE)
```

With this set of permissions, all IDs in the GRP class are authorized to use ICSF facilities and the shared SYSTOK-SESSION-ONLY token.

Encoding HTML Tags in Data

You can set the HTMLENCODE parameter to control whether HTML tags in data are encoded as plain text or HTML tags. The value ON encodes the tags as plain text so that the browser does not consider them to contain executable code. This prevents an attack on the server by inserting executable code into data. You can set the value of HTMLENCODE to ON in the server Miscellaneous Settings, available from the Workspace page of the Reporting Server browser interface. Doing so places this setting in the edasprof.prf file and enforces it server wide.

Note: You can store the value in the edaserve.cfg file using the *Miscellaneous Settings* option of the *Workspace SET* button on the Workspace ribbon, as described in [Miscellaneous Settings](#) on page 622.

Access to Connection Information in TIBCO WebFOCUS Procedures

The TIBCO WebFOCUS Client can pass a browser cookie or HTTP header to the Reporting Server for use in WebFOCUS procedures.

The following server built-in functions can be used to retrieve this information and use it in a procedure.

- The GETCOOKI function can be used to retrieve the browser cookie value.
- The GETHEADR function can be used to retrieve variable values from the HTTP header.

For more information about WebFOCUS functions, see the *TIBCO WebFOCUS® Using Functions* manual.

In addition, three Dialogue Manager system variables contain connection information.

- &FOCSECUSER returns the user ID of the connected user,
- &FOCSECGROUP returns the primary group ID for the connected user, if there is one.
- &FOCSECGROUPS returns the list of group IDs for the connected user.

Note: The length limit of the returned list is 4000 characters.

Configuring Multi-Tenant Deployment

In multi-tenant software architecture, each tenant (client) shares certain computing resources and application code with other tenants. In addition, each client has its own individual resources that are logically isolated from those belonging to other tenants. Metadata associates each resource with the correct tenant.

Within the server framework, in a typical setup, each tenant has its own:

- Application folders.
- Groups and users.
- DBMS connections

Configuring Security Templates for Multi-Tenant Environments

Very often, multiple tenants in a multi-tenant environment have an identical security structure. Each tenant has an identical set of application folders and groups that conform to a standard naming convention that includes the tenant ID in the name of each subgroup and application folder. The folder and group setup is repeated from one tenant to the others. In most cases, tenants have no access rights outside of their main application folder.

In this type of configuration, you can establish file privileges for a model tenant and its associated groups, and dynamically apply them to other tenant groups by creating a template based on the model tenant. The template defines regular expressions that identify which actual group names should be assigned file privileges based on that template. Each group ID should match only one template.

Procedure: How to Configure Security Templates

The following steps outline the tenant management process in the Reporting Server.

1. Identify the tenant root application folder.


The tenant root application folder can be under approot or can be a mapped application or SQL mapped folder. The application must be in the APP PATH. For example, c:\ibi\apps\tenants.

2. Set the permission for this folder to no access for all non-server administrator roles. in order to prevent tenants from accessing folders assigned to other tenants.
3. Using the established naming convention, create the model tenant folder under the tenant root. Give the model tenant groups explicit access permissions to the model folder and its subfolders.

The model tenant groups must follow the tenant group structure and naming conventions established in the security software for the installation. The model group must be registered in the admin.cfg file. If the security provider is PTH, the tenant groups must also be registered in the admin.cfg file.

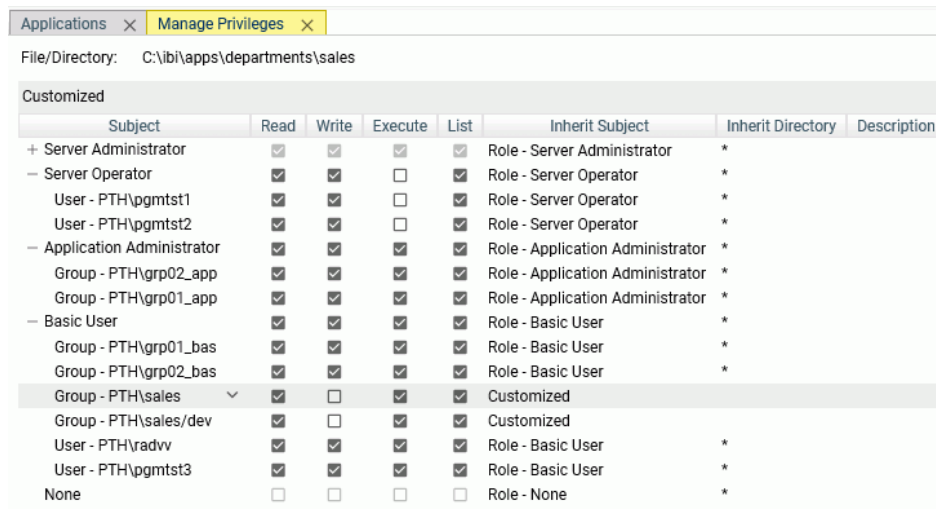
For example, using the departments application as the tenant root, create the *sales* application under the tenant root. The server administrator creates the groups *sales*, *sales/dev*, *sales/basic*, and any other groups needed.

In the following image, the tenant root folder is called *departments*, and the *sales* application is the model application. The sales application is not only a model, it is an existing group or tenant that must be registered, with privileges to the application *departments/sales*. Once you create one registration you can assume it as the model and create a template based on it:

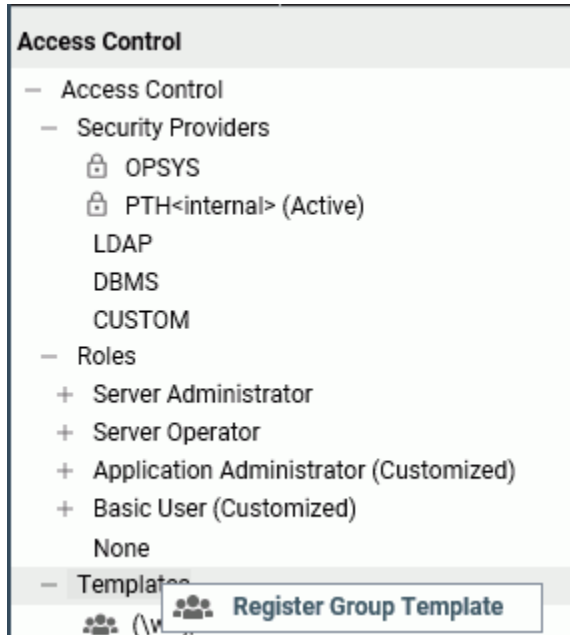


The server administrator defines the file and directory privileges for the sales, sales/dev and sales/basic groups on the sales folder and its children. These privileges will be replicated in groups for other tenants when their group ID matches one of the established template group IDs.

The following image shows that the *sales* group, whose role is Basic User, has been given Read, Execute, and List privileges on the sales folder:



4. Establish a template that specifies how the model access rules are replicated to other tenants.
 - a. On the Access Control page, right-click the Templates folder and select *Register Group Template* from the context menu:



The Group Template Registration page opens:

- b. Enter the following values to establish the rules for deriving access control privileges for groups that match the template, and click *Register*:

Template Group ID

Is a regular expression that must contain a tagged expression. The tagged expression is used to identify a part of the model group ID that will be replaced by the corresponding characters in the actual group ID passed to the Reporting Server for the connected user.

For information about regular expressions, see [Summary of Regular Expressions](#) on page 111.

Model Group

Identifies the group that will be replicated.

Exclude Group IDs

Is a regular expression that identifies any group IDs that should be excluded in the template matching process.

Replace Pattern

Is used for `admin_privilege` strings to model file/directory privileges from the model group registration to a target group.

For example, the sales application has the `sales/dev` group that is registered with specific privileges on the sales folders. When the finance application is created, it will have a group named `finance/dev` that will be assigned the same privileges on the finance folders that the `sales/dev` group has on the sales folders. The finance groups will not be registered, but the template describes how to replace the word `sales` with the word `finance` when the group name ends with the characters `/dev`.

When you register the template, the following attributes are added to the `admin.cfg` file:

```
admin_group_template = (/w+)/dev
BEGIN
  model_group = sales/dev
  file_replace_pattern = departments\\(sales)
END
```

For example, the template group ID pattern `(/w+)/dev` specifies that any word followed by the characters `/dev` matches the template pattern.

The replace pattern `departments\\(sales)` indicates that the text `sales` will be replaced. If the group ID `finance/dev` is passed to the Reporting Server, the tagged expression in the Template Group ID will cause the characters `finance` to replace the characters `sales` in order to determine the file privileges for the connected user. Therefore, a user in the group `finance/dev` will inherit the privileges under the `finance` tenant application that the group `sales/dev` has under the model sales application.

To create a template that applies to the group *sales*, you need to make sure that you exclude other group names that start with the characters *sales/* because without the exclusion, *sales* will match all of the templates, and it may be assigned the wrong one. The following image shows the template for the sales group:

Activate Providers x Group Template Registration x

Example:
Here are the steps to model groups Accounting/Developers, Finance/Developers or HR/Developers on a registered group Sales/Developers, and repl: Sales/Developers with tenant\accounting, tenant\finance or tenant\hr.

- Set 'Template Group ID' to regular expression '(w+)/Developers'
- Set 'Model Group' to Sales/Developers
- Set 'Replace Pattern' to regular expression 'tenant\\(sales)'

[click here for the summary of regular expression](#)

? Template Group ID (w+) example: (w+)/Developers

? Model Group PTH\sales

? Exclude Group IDs (/) example: Support or Operators|Marketing

? Replace Pattern departments\\(sales) example: tenant\\(sales)

Update Cancel

The regular expression `(/)` in the Exclude Group IDs field specifies that group names that contain a slash do not match the template. For information about regular expressions, see [Summary of Regular Expressions](#) on page 111.

At this point, the system is ready to add new tenants.

To add a new tenant (for example, *finance*), the server administrator must create the finance application under the tenants root application. The installation must provision groups and users (that conform to the established naming conventions) in the relevant security provider, for example, LDAP, OPSYS, or Custom. This is done outside of TIBCO software. The new tenant is now ready to be used. No access control for it needs to be established by the server administrator. The server will dynamically assign the access control privileges based on the model and the template.

For example, the user *fbas1* (which is part of the finance group) logs on. This user has Read, Execute, and List privileges on the finance folder because the sales model group was given those privileges on the sales folder.

The sales tenant is not visible to this user.

Reference: Summary of Regular Expressions

A regular expression is a string containing a tagged expression. A tagged expression is a portion of the string that will be replaced with actual characters passed to the Reporting Server in the group ID for a user.

Tagged expressions must be enclosed in parentheses. The backslash character (\) is a special character in tagged expressions. If a group ID actually contains a backslash character, you indicate that you want the backslash character to be treated as a normal character by entering two successive backslash characters (\\).

The following can be used to create the tagged expression:

`abc`

Matches *abc* anywhere within the string.

`(abc)`

Matches *abc* anywhere within the string, but the parentheses act as a tag.

`[]`

Defines a character class (or set) that matches any one of the characters in the class. For example, `[abc]` matches the character a or b or c. The expression `[x-y]` matches any character from x to y.

`.`

Matches any single character except newline.

`\w`

Matches any word character string (alphanumeric plus "_").

`\W`

Matches any non-word character.

`\s`

Matches any whitespace character.

`\S`

Matches any non-whitespace character.

`\d`

Matches any digit.

`\D`

Matches any non-digit character.

`\t`

Matches a tab character.

`\r`

Matches a return character.

`\f`

Matches a formfeed character.

`\e`

Matches an escape character.

`\b`

Matches a word boundary or a backspace. For example, `test\b` matches `test`, but not `testing`. However, `\b` matches a backspace character when specified inside a class (that is, `[\b]`).

`\B`

Requires that there is no word boundary. `hello\B` matches *hello*, but not *hello there*.

`^`

Matches characters only at the beginning of the string. For example, `^abc` matches `abc` at the beginning of the string.

`$`

Matches characters only at the end of the string. For example, `abc$` matches `abc` at the end of the string.

`|`

Specifies alternative matching characters. For example, `a|b` matches either `a` or `b`. This metacharacter can also be used with words, for example, `abc|def`.

`[^abc]`

Matches a character that is not in the set. `[^abc]+` will match such strings as *hello*, *test*, and *perl*.

`\`

Is the escape character. For example, `*` matches the `*` character. Use the backslash (`\`) character to escape (remove the special meaning of) characters that have significance in a regular expression.

`(?i)`

Ignores case. For example, `(a(?i)b|c)` matches `aB`, `c`, and `C`.

You can follow any character, wildcard, or series of characters and/or wildcards with a repetition indicator:

*

Matches 0 or more occurrences of the character sequence.

+

Matches 1 or 0 occurrences of the character sequence.

?

Matches 0 or more occurrences of the character sequence or the shortest match.

{ }

Is the repetition modifier.

{ *n* }

Matches exactly *n* occurrences of the character sequence.

{ *n*, }

Matches at least *n* occurrences of the character sequence.

{ *n*, *m* }

Matches at least *n* but not more than *m* occurrences of the character sequence.

Configuring Multi-Tenant Application Folders

The main application folder for each tenant should contain its unique metadata (synonyms), uploaded data files and HOLD files, StyleSheet files, and server-based procedures. The tenant should have no access to other tenant folders. However, a tenant may have access to common installation-wide folders under `aproot`, which is useful when metadata and utilities are shared between tenants. Nested folders under the main tenant application folder can be added as needed.

Configuring Multi-Tenant Groups and Users

Each tenant will have at least two groups. One group (registered to the application administrator role) has read/write access to the tenant folders and is able to create metadata and other files. The other group (registered to the basic role) can have only execute and list privileges on the tenant folder.

More groups can be added as needed.

It is recommended that groups and users be defined externally (in LDAP, for example) so that user provisioning, password expiration, and other administrative tasks can be done with the standard tools. However, the PTH security provider with server-maintained users and groups can be used as well.

Configuring Multi-Tenant DBMS Connections

DBMS connections can be defined for all tenants in the `edasprof` file or in the group profile for each tenant.

Configuring Security Providers for Multi-Tenant Deployment

The server administrator can give each tenant administrator the ability to add and remove users independently by configuring a separate security provider for each tenant. Appropriate LDAP server privileges can be given to selected users for each tenant. Alternatively, a single provider can be used. These setup choices are part of the LDAP configuration techniques for the site.

Naming Conventions for Multi-Tenant Deployment

A naming convention should be developed for creating application folders and groups that contain the tenant ID as a suffix or prefix. In the following example, the tenants have IDs 01 and 02, and folders and groups are named the following:

```
app01, app02, grp01_app, grp01_bas, grp02_app, grp02_bas
```

This will make administrative tasks more transparent and allow easier replication of the configuration when new tenants are created.

Configuring Application File Privileges for Multi-Tenant Deployment

It is recommended that an application named *tenants* be created under `aproot` and that all application files be created under the tenants application.

Once the tenants application is created, the administrator should remove all privileges to it from the *application administrator* and *basic* roles. This will ensure that tenants do not have any access to the applications owned by other tenants:

File/Directory: C:\jbi\apps\tenants

Customized							
Subject	Read	Write	Execute	List	Inherit Subject	Inherit Directory	Description
+ Server Administrator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Administrator	*	
- Server Operator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Operator	*	
User - PTH\pgmtst1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Role - Server Operator	*	
- Application Administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Customized		
User - PTH\pgmtst2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Application Administrator	C:\jbi\apps\tenants	
Group - PTH\grp02_app	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Application Administrator	C:\jbi\apps\tenants	
Group - PTH\grp01_app	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Application Administrator	C:\jbi\apps\tenants	
- Basic User	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Customized		
Group - PTH\grp01_bas	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Basic User	C:\jbi\apps\tenants	
Group - PTH\grp02_bas	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Basic User	C:\jbi\apps\tenants	
Group - PTH\sales	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Basic User	C:\jbi\apps\tenants	
Group - PTH\sales/dev	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Basic User	C:\jbi\apps\tenants	
User - PTH\radv	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Customized		
User - PTH\pgmtst3	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - Basic User	C:\jbi\apps\tenants	
None	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Role - None	*	

Next, give each tenant group access to its own application. For example:

Grp01_app has all file privileges registered to tenants/app01.

Grp01_bas has only execute/list file privileges registered to tenants/app01. Both tenant 01 groups have *no* access outside of their folder.

Effect of Multi-Tenant Deployment on File Upload and HOLD Files

Uploading can be done to a designated subfolder of the tenant application. Even basic role users need write access to the upload folder.

For example:

Grp002_bas has all file privileges registered to tenants/app01/upload

Alternatively, each user can use his home folder and not share the files, even with users of the same tenant.

Shared Metadata Under Multi-Tenant Deployment

The administrator may choose to share all synonyms and procedures for all tenants. A common application folder outside of the tenant application should be created with only list/execute permissions for all tenant groups. This folder is updated centrally for the whole installation. In this case, the tenant folders can be used only for upload and HOLD files.

Chapter 3

Managing Applications

An application is a platform-independent repository for a group of related components, such as procedures, Master and Access Files, data files, HTML files, PDF files, and image files.

You can use a variety of application (APP) commands to control the application environment, including the application itself, its component files, and its search paths.

In this chapter:

- [What Is an Application?](#)
- [Procedures and Metadata on the Application Tree](#)
- [Managing Applications and Paths](#)
- [Indexing Applications For Global Searches](#)
- [Application Commands Overview](#)
- [Search Path Management Commands](#)
- [Application and File Management Commands](#)
- [Output Redirection Commands](#)
- [Application Metadata Commands and Catalog Metadata](#)
- [APP HELP](#)
- [Restricting the Use of APP Commands](#)
- [Accessing Metadata and Procedures](#)
- [Allocating Temporary Files](#)
- [Temporary Space Usage and Location](#)
- [Temporary Disk Space Usage for Non-PDS Deployment](#)
- [Application Tools](#)

What Is an Application?

An application is a platform-independent repository for a group of related components, such as procedures, Master and Access Files, data files, HTML files, PDF files, and image files. It provides an area on the server that both confers a unique identity on the application components and facilitates the sharing of components across applications in an organized manner. This construct also simplifies the process of moving a user application from one platform to another and of deploying PC-developed applications.

These components are physically grouped together on an application-by-application basis for run-time execution. This physical grouping can be within an application under a common root or a mapping to an application anywhere in the file system. The physical application or mapped name is referred to as the application name in this document. A comprehensive set of application (APP) commands are provided to control/manipulate the application components, as well as to facilitate applications that can be written and deployed to any platform.

The physical location of an application and its components is determined by a configuration parameter called `aproot`. This parameter is set at installation time and stored in the server configuration file, `edaserve.cfg`. The default value is dependent on the platform, relative to the install ID home directory, where applicable.

Application directories can be nested. A nested application directory is an application created within a higher-level application. For more information, see [Nested Application Directories](#) on page 150.

You can also create a home application directory for each user. Providing a user home application gives each user a place where he has full control to create, change, and run his applications. For more information, see [Home Application Directories for Users](#) on page 153.

The various operating systems on which the product runs have their own behaviors for physical files and how directories or components are referenced. In addition, some are case-sensitive and some are not. For example, on Windows, the files `abc` and `ABC` refer to the same file, regardless of how it is stored on disk (`aBc`, for example). However, on UNIX they are all different files. Additionally, z/OS under PDS deployment makes file names uppercase when they are actually saved but they can generally also be referenced in lowercase or mixed case after they have been created. Some operating systems also allow spaces in file names and some do not. To co-exist within the various platforms on which the product runs, the product rules are to either always use APP commands or product tools such as the Reporting Server browser interface to create apps for application files (which will create them in the appropriate case), or to use lowercase names with external tools (such as `mkdir myapp` and `vi mytest.fex`), as they will save appropriately and work within the APP framework. Additionally, spaces in file names are not allowed. The use of external tools is also discouraged, as they may not act in the same way as internal tools.

Note:

- ❑ Where directories are referenced in the chart, lowercase application directories are created below the `aproot` value. Actual application names must be lowercase and must not contain spaces. For z/OS PDS deployment, data sets are created for each component type using the `aproot` value as the high-level qualifier.

- ❑ Some server features will not work without nested applications enabled. Among those are home application and file upload.

Platform	Default Value for <code>aproot</code>
z/OS PDS Deployment	<code>aprootvalue.appname.component_type</code>
z/OS HFS Deployment	<code>.../ibi/apps</code>
UNIX	<code>.../ibi/apps</code>
Linux	<code>.../ibi/apps</code>
Windows	<code>...\\ibi\\apps</code>
IBM i (formerly known as i5/OS)	<code>.../ibi/apps</code>

Two applications are provided during installation: a default application called *baseapp* and an application in which you can generate legacy sample files called *ibisamp*. In addition, when you connect to the server, a temporary directory called *foccache* is added as the first directory in the search path. When you want to be able to reuse data within the same browser session, you can store the data in the form of a HOLD, SAVE, or SAVB file in the *foccache* directory. As long as the browser session remains active, the files stored in the *foccache* directory can be referenced in requests.

Access to a particular application component can be explicit or implicit. Implicit access is dependent upon the search path in effect at the time of execution. The search path always includes the default application, *baseapp*. There is no need to explicitly declare this application.

You can change the search path from the Reporting Server browser interface, from the Data Migrator desktop interface, or from within your application code. You can also change the search path temporarily to add application names to the beginning or end of an existing search path. APP commands are described briefly in [Application Commands Overview](#) on page 174 and in detail later in this chapter.

In addition to explicit APP names under APPROOT or an APP MAP command, which might look like *myapp/myproc*, there are also special reference names for internal locations that may be useful. They are:

Reference Name	Description
_edatemp	Current directory location. For server use, this is the EDACONF edatemp agent directory (that is, ts000001, where 000001 is tscomid of the agent). For PDS Deployment, this is a temporary HFS location. For edastart -t, -x and -f uses, this is the current user directory.
_edahome	EDAHOME installation location.
_edacnf	EDACONF configuration location.

While the EDAHOME and EDACONF catalog locations are internally already on a server search path, special handles such as these allow you to explicitly reference names of the form `_edahome/catalog/sysapps`, which is one of a number of internal catalog tables.

Note: For platforms that support the Universal Naming Convention (UNC), you can specify a network drive for approot. The UNC must be:

- Minimally one folder below the initial shared location.
- Not contain spaces. For example:

```
\\mynode\myshare\accting
```

Generating Samples and Tutorials

You can generate several types of sample files for WebFOCUS and TIBCO® Data Migrator and save them in an application in order to run sample procedures and test features and configurations.

The Reporting Server browser interface Applications page has an option on the *New* menu named *Tutorials*, which you can click to open the *Create Tutorial Framework* page. You can also right-click an application folder, click *New*, and then *Tutorials*, from the context menu.

Following is the list of available tutorials.

- WebFOCUS - Getting Started Demo

- WebFOCUS - Retail Demo (can include Rserve demo if Rserve is configured)
- WebFOCUS - SSAS Cube Join Demo
- WebFOCUS - State Population Demo
- WebFOCUS - Custom SQL Security Provider
- Data Migrator - General
- Data Migrator - Iterator
- Data Migrator - File Listener
- Data Migrator - Star Schema
- Create Legacy Sample Tables and Files

In the past, a static application directory named `ibisamp` was created and pre-populated with a number of classic sample/demo files from past releases. The `ibisamp` folder is still created as part of the base installation process, but it is no longer automatically populated.

The current tutorial options have a wider variety of demos, but only load when selected. The classic sample/demo files are still available by choosing *Create Legacy Sample Tables and Files* from the Tutorial drop-down list.

Because the change is implemented at the folder level, you can choose to continue using `ibisamp` as the location for creating the legacy files, or you can select any other folder. The drop-down option for Legacy samples creates most, but not all, of the sample files that used to be in `ibisamp`. Some of the prior files were actually Data Migrator-related, and those have been moved to their own tutorials.

Under z/OS PDS Deployment there are some additional restrictions. The z/OS PDS JSCOM Listener option must have been selected at installation time. The PDS JSCOM option also co-installs some HFS source files that are used by the various create options. If the HFS files for a z/OS PDS Deployment are not present, request to create a tutorial will produce a specific *HFS files not found* message. Note that not all tutorials are implemented for z/OS, and those will produce a *not implemented* message, if selected.












Procedures and Metadata on the Application Tree













The Applications page groups all application files, including procedures, synonyms, HTML files, data flows, user functions, and other files on a single application tree. By default, files in an application are listed in order of file type. The file type is indicated by a unique icon for each type that displays to the left of the file name in the tree. However, you can sort or filter the list to generate custom views of each application.









Right-clicking an item brings up a context menu. The items you see on the context menu depend on your role and the type of file. Some items are only available to administrators or other users with administrator privileges. Some of the items on the context menu provide shortcuts to items on the *Workspace* or *My Console* menu.

Reference: Application Tree Icons

The following table lists the icons displayed for each type of file on the application tree.

Icon	Type of File
	Closed application directory
	Opened application directory
Metadata	
	Synonym (.mas)
	Cluster synonym (.mas)
	Business View synonym (.mas)
Procedures	
	Stored Procedure (.fex)
	Data Migrator Flow (.fex)
	Data Migrator Flow with Change Data Capture (CDC) (.fex)
	DBMS SQL Flow (.fex)
	Direct Load Flow (.fex)
	Direct Load Flow with Change Data Capture (CDC) (.fex)

Icon	Type of File
	User Function (.fex)
	Scheduled Only (.fex)
Documents	
	HTML file (.htm, .html, .shtml, .htt, .mht, .mhtml, .cfm, .tpl, .hta, .htb)
	Microsoft Excel® Document (.xls, .xlsx, .xlsb, .xht, .xltx, .xlsm, .xltm)
	Adobe Acrobat Document (.pdf, .ai)
	XML Document (.xml, .wsd, .xsd, .wsdl, .mxml, .gcl, .xul, .dtd, .xsl, .xslt, .axl)
	JavaScript Object Notation (.json)
	Microsoft PowerPoint® Document (.ppt, .pptx, .pptm)
	Microsoft Word® Document (.doc, .docx, .docm, .dot, .dotx, .dotm)
Other	
	SQL Script (.sql)
	Data File (.foc, .ftm, .dat, .txt, .csv, .tab, .bdat, .data, .tmp)
	Graphics File (.jpg, .jpeg, .jpe, .svg, .jfif, .tif, .tiff, .ico, .gif, .bmp, .png) Note: The text inside the icon identifies the type of image file.

Icon	Type of File
	Style File (.sty, .focstyle, .css)
	Archive File (.zip, .rar, .tar, .jar, .war)
	Script File (.js, .jcs, .omi, .cbl, .hti, .jji, .vbs, .sh, .ctl, .bat, .t3i, .jcl, .ps)
	Log/Trace File (.log, .trc, .hto, .t3o, .sta, .msg)
	Configuration File (.cfg, .ini, .prf, .err, .nls)
	Language File (.lng)
	Other Document (.prn, .dif, .fmu, .wp, .ifp, .lzx, .syl, .dmc, .tdl, .bst, .adr, .eps, .swf, .as, .trf, .idx)
	MAINTAIN File (.mpt, .mnt, .fcm, .wfm, .wri, .wxi)

Reference: Context Menu Options for Application Directories

The following options are available on the Applications page ribbon:

New. Enables you to create new application files. The available context menu options are:

- Cluster Business View.** Opens the Synonym Editor.
- Flow.** Opens a flow page.
- Advanced/Custom Copy.** Starts the custom copy process.
- Procedure.** Opens the New Procedure text editor.
- Text File.** Opens a text box in which you can create a new file as text.
- Upload Raw Data.** Enables you to upload image or data files to an application directory. No synonym is created for data files.

- Application Directory.** Enables you to create a new application directory folder.
- Tutorials.** Opens the Create Tutorial Framework page where you can create samples tables and metadata.

Get Data. Opens the Get Data page for uploading files or configuring adapters and connections.

Filter. Enables you to customize the items that display on the application tree. If there are nested application directories and you want to see the selected files from all of them in a single list, click the ellipsis button (...) on the files pane. The available context menu options are:

- Procedures.** The options are *All Procedures* or *Scheduled Only*.
- Synonyms.** The options are *All Synonyms*, *Cluster Only*, or *Business View Only*.
- Advanced.** Opens the Filter Applications Tree page, where you can specify file properties, select among all of the types of files available, and select file statistics such as size or date modified. For more information, see [How to Filter Items on the Application Tree](#) on page 163.
- Impact Analysis.** Reports on your use of synonyms, columns, and procedures. The choices are Synonyms by Procedure, Procedures by Synonym, Columns by Procedure, and Procedures by Column. In addition, you can select a Summary or Detailed flow report.
- Recent.** Opens the list of files the user opened recently (history.txt in the myhome application for the user).
- Favorites.** Opens the list of files the user added as favorites.

Manage. Enables you to manage your applications. The following context menu options are available:

- Settings.** Opens the Application Settings pane.
- External Repository.** Has the following options.
 - SQL Repository.** Provides options to create and manage SQL repositories, allowing you to store the contents of an application folder in an SQL database.
 - WebFOCUS Client Repository.** Provides options to link to or list and delete a WebFOCUS Client Repository.
- Application Path.** Opens the Application Path Configuration page under the My Console menu while leaving the Application Directories tree open in the left pane.

- Schedule/Email.** Has the options *Scheduler Agents*, *Scheduled Events*, *Log and Statistics*, and *Force Scheduler Scan*. Enable you to view various scheduler pages and logs and statistics while leaving the Application tree open in the left pane.
- Create My Home.** Creates the home application directory of the user.
- Shared Resources.** Opens a pane that enables the user to share personal files and folders under their myhome directory and shows the shared resources the user has configured.
- Preferences.** Opens the Application Preferences page, where you can customize the tree display. The available options are:
 - Show edahome, edaconf, edaprfu, scaroot, edatemp, edalog and foccache of all users.** Shows internal directories on the application directories tree for users with Server Administrator privileges.
 - Show Descriptions on Applications Tree.** Shows the file description on the files pane.
 - Show Also Applications not in APPATH.** Adds an Inactive Directories tree listing all of the applications not in the path to the Application Tree. When the Inactive Directories tree is displayed, the context menu option changes to Hide Applications not in Path.

Note that this option is only available for the server administrator and users with server administrator privileges. Users with these privileges can manage (add, delete, copy, move, or modify files) directories that are not in the Application Path.
 - Swap File/Dir Description and Name.** Shows the description, if any, where the file name displays in the default layout.
 - Show Files from Nested Applications.** Shows files from nested applications in the list of files. The default display shows nested application directory names, and you can double-click them to show the files within them.
 - Show Application Profile.** Shows the current application profile.
 - Show Files on Application Tree.** Enables you to open the list of files directly in the Application Tree instead of in a separate pane.
- Reset.** Resets the application page to its default settings and enables the user to show or hide the different panes.

Reference: Context Menu Options for an Application Directory

Like the top-level of the Application Directories tree, application directory folders include New, Schedule and E-mail, and Impact Analysis options, as described in [Context Menu Options for Application Directories](#) on page 124.

Application directory folders also include the following right-click options:

- Refresh.** Refreshes the list of files in the application.
- New.** Enables you to create a new cluster Business View, synonym, custom copy, data flow, procedure, text file, upload raw data, application directory, or tutorials.
- Quick Copy.** Enables you to copy data when you have configured SQL adapters.
- Schedule and E-Mail.** Opens the Scheduler Agents or Scheduled Events page for the application or forces a Scheduler scan.
- Create Index.** Allows you to create an index for the selected mapped application directory.
- Update Index.** Allows you to update an existing index for the selected mapped application directory.
- Delete Index.** Allows you to delete an existing index for the selected mapped application directory.
- Log and Statistics.** Opens the Log and Statistics page where you can view a range of log and statistics reports.
- Impact Analysis.** Allows you to run impact analysis and flow reports for the application.
- Copy.** Copies the application directory.
- Delete.** Deletes the application directory and is only available for application directories that you created.
- Cut.** Cuts the application directory and is only available for application directories you created.
- Paste.** Pastes the application directory.
- Rename.** Opens a dialog box for renaming the application directory.
- Privileges.** Opens the Manage Privileges pane for this application directory. This option is only available to Server Administrators.

- ❑ **Properties.** Opens the Properties pane that shows the location of the application directory, number of files in the directory, the date the directory was last modified, and the description. You can also edit the description using the text box provided.

Reference: Context Menu Options for All Files in an Application Directory

All files include the standard right-click options of Cut, Copy, Paste, Rename, and Delete.

You can multi-select files before clicking an option on the right-click menu. To select a group of contiguous files, click the first file, hold down the Shift key and click the last file. To select discontinuous files, hold down the Ctrl key while clicking the files you want to select.

Files also include the following right-click options:

- ❑ **Open.** Opens the file in the editor associated with that type of file (not available for Quick Query, HTML, and Synonym).
- ❑ **Properties.** Opens the Properties pane that shows the location of the file, the date it was last modified, the Job Type, the description, and the privileges associated with the file.
- ❑ **Download.** When this option is selected, the file is transferred to the Downloads folder of the user. On Windows, this folder is, by default, the following directory:

`C:\users\username\Downloads`

If you right-click a synonym, the shortcut menu gives you a choice of Master File or Access File.

- ❑ **Copy.** Copies the application file. Multi-select is supported. You can hold the Ctrl key to add files.
- ❑ **Delete.** Deletes the application file and is only available for application directories that you created.
- ❑ **Cut.** Cuts the application file and is only available for application directories you created.
- ❑ **Rename.** Opens a dialog box for renaming the application file.
- ❑ **Add to Favorites.** Adds the file to your list of favorites. You can view the list of favorites by clicking *Filter* on the ribbon and selecting *Favorites*.
- ❑ **Privileges.** Opens the Manage Privileges pane for this application file. This option is only available to Server Administrators.

- Properties.** Opens the Properties pane that shows the location of the application file, the date the file was last modified, the description, and the privileges for the connected user. You can also edit the description using the text box provided.

Reference: **Context Menu Options for Stored Procedures**

Stored procedures include the right-click options Open, Delete, Cut, Rename, and Properties, as described in [Context Menu Options for Application Directories](#) on page 124.

Stored procedures also include the following right-click options:

- Open Advanced.** Attempts to open the procedure in the Flow Editor.
- Run.** Runs the file if it is an executable file.
- Run Advanced.** Opens a menu containing the following options:
 - Submit with E-mail.** Enables you to manage email notification when the procedure starts and/or completes
 - Run Changing Defaults.** Allows you to enter values for variables and then runs the file if it is an executable file.
 - Run As.** Prompts for credentials prior to running the procedure if the `connections_autoprompt` profile setting is `y` and the *Auto Prompt Connection Credentials* setting is `y` for the connection being used.
 - Run Stress.** Opens a pane that enables you to set the number of threads, the interval, and the time to keep alive the procedure when it runs. You can also choose to generate statistics and comparisons for this run.
 - Run Stress Deferred.** Opens a pane that enables you to set the number of threads, the interval, and the time to keep alive the procedure when it runs deferred. You can also choose to generate statistics and comparisons for this run.
 - Debug.** Enable you to run the procedure in debug mode. You can step through the code, viewing the result for each line. For more information, see [Debugging a Stored Procedure](#) on page 473.
- Schedule and E-Mail.** Enables you to manage email notification when the procedure starts and/or completes, and to change default values of amper variables. It also enables you to open the list of Scheduler Agents and Scheduled Events.

- Logs.** Provides the following options:
 - Last Log.** Displays the log from the last time the procedure was run.
 - Last Output.** Displays the output from the last time the procedure was run.
 - Log and Statistics.** Opens the Logs and Statistics page where you can view a range of log and statistics reports.
- Analysis Reports.** Provides the following options:
 - Impact Analysis.** Reports where this procedure is found and its properties.
 - Dependencies Analysis.** Displays the Dependencies Analysis page. It enables you to see the synonym(s) and other procedures that are used by the procedure.
 - Data Lineage.** Reports on the fields in the procedure and their properties.

Reference: **Context Menu Options for TIBCO Data Migrator and Direct Load Flows**

Data Migrator and Direct Load Flows include the Schedule and E-mail and Impact Analysis options, as described in [Context Menu Options for Application Directories](#) on page 124. They include the Run, Run Advanced, Logs, and Dependencies Analysis options, as described in [Context Menu Options for Stored Procedures](#) on page 129.

Data Migrator and Direct Load Flows also include the following options:

- Open Flow.** Opens the flow in the Data Flow Editor.
- Submit.** Submits the flow for processing.
- Flow Report.** Displays details about the flow, including source and target information, load options, SQL select statements, transformations, execution properties, and record logging.

Reference: **Context Menu Options for Quick Queries**

Quick Queries include the Schedule and E-mail and Impact Analysis options, as described in [Context Menu Options for Application Directories](#) on page 124. They also include the Run, Run Advanced, Submit with Options, and Logs options, as described [Context Menu Options for Stored Procedures](#) on page 129.

Reference: **Context Menu Options for User Functions**

In addition to the standard file options (Cut, Copy, Paste, Delete, and Properties), User Functions include the Test right-click option, which tests the function.

Reference: Context Menu Options for HTML Files

In addition to the standard file options (Cut, Copy, Paste, Delete, and Properties), HTML files include the Run and Logs options, as described in [Context Menu Options for Synonyms](#) on page 131.

HTML files also include the following options:

- Edit.** Enables you to edit the HTML code.
- View.** Enables you to view how the file displays in a browser.

Reference: Context Menu Options for Synonyms

Synonyms include the Analysis Reports option, as described in [Context Menu Options for an Application Directory](#) on page 127. They also include the Download, Copy, Delete, Rename, and Properties options, as described [Context Menu Options for Stored Procedures](#) on page 129.

Synonyms also include the following options:

- Open.** Opens the synonym in Data Assist.
- Sample Data.** Runs a sample report against a synonym.
- Data Profiling.** Runs a report against a synonym listing its segments, fields, and field information, with a drill-down report to field values.
- Analysis Reports.** You can select Impact Analysis, which reports where this synonym is found and its properties, or Dependencies Analysis, which displays the Dependencies Analysis page. It enables you to see the applications and synonyms that contain or have references to this synonym.
- Metadata Management.** Provides options to create a cluster synonym, to refresh (update) synonyms (only available for base synonyms, not cluster synonyms), to edit the Master File in the text editor, to open the Access File in the text editor, and to prepare translation files for the metadata in the synonym.
- Data Management.** Enables you to recreate a DBMS table, drop a table, insert sample data, delete all data, and show/modify data.
- Quick Copy.** Opens the Quick Copy page where you can specify the target parameters for the new synonym and table.
- Custom Copy.** Opens Data Assist for editing and enhancing the synonym.

- Flow.** Opens a demo flow page.

Managing Applications and Paths

Path management tasks are available from the Applications page, which is accessed on the *Applications* page. The Application Directories tree displays on the resources pane. A set of buttons provide the quickest way to initiate path management tasks.

Creating and Mapping Applications

Applications are designed to group related components.

- When you create a new physical application, it becomes available for addition to the search path in a selected profile.
- When you map an application name to an existing physical location outside of approot, the mapped application becomes available for inclusion in the search path of a selected profile.

Applications can be created and mapped in either the Reporting Server browser interface or the Data Migrator desktop interface.

Application directory names must comply with the following rules:

- The maximum length is 64 characters.
- For multi-part Application directories, the total length is limited to 512 characters.

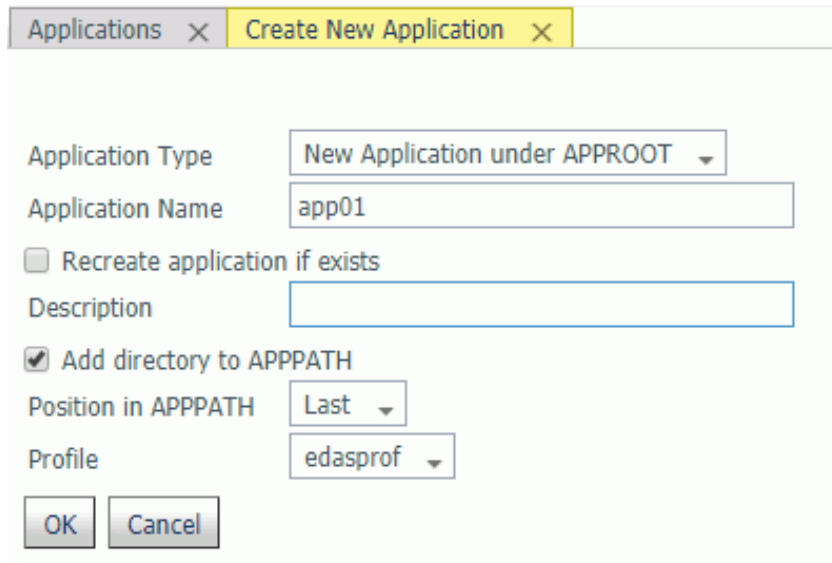
❑ Names can contain any valid alphanumeric character except the following:

' '	space character	,	comma
*	star character	.	dot
?	question mark	:	colon
<	less than sign	"	double quotation mark
>	greater than sign	'	single quotation mark
&	ampersand	\t	tab
\	backslash	\0	null terminator character
	vertical bar	/	slash
=	equal sign		

***Procedure:* How to Create an Application**

1. On the Reporting Server browser interface, go to the *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Either click *New*, then *Application Directory* on the Reporting Server browser interface ribbon, or right-click an application and click *New*, then *Application Directory*.

The Create New Application page opens.



Applications × Create New Application ×

Application Type New Application under APPROOT ▾

Application Name app01

Recreate application if exists

Description

Add directory to APPPATH

Position in APPPATH Last ▾

Profile edasprof ▾

OK Cancel

3. Use the default *Application Type*, *New Application under APPROOT*.
4. Enter a name in the *Application Name* field.
5. Optionally, select the *Recreate application if exists* check box.

Warning: Choosing this option will overwrite the existing application and any content in it.

6. Enter a description in the *Description* field.
7. The *Add directory to APPPATH* option is the default. Optionally, you can decide not to add the directory.
8. Select a position from the *Position in APPPATH* drop-down menu. The options are *Last* and *First*. The default is *Last*.
9. Select a profile from the *Profile* drop-down menu. For Server Administrators, the default is *edasprof*. For all other users, the default is the user profile.
10. Click *OK*.

The application is added to the *Application Directories* folder.

Procedure: How to Map an Application to a Physical Directory

You can map an application name to a physical directory anywhere in the file system. This application name can be then used in APP commands.

Application mappings can be added and deleted based on profiles from either the Reporting Server browser interface or the Data Migrator desktop interface.

1. On the *Applications* page, either click *New*, then *Application Directory* on the ribbon, or right-click an application, select *New*, and then *Application Directory*.

The Create New Application page opens.

2. Select *Application Mapping to Disk* or *SQL Repository* from the *Application Type* drop-down menu.

3. Enter a name in the *Application Name* field.
4. Accept the default *Physical location*, enter a different location, or click the selector button (...) and navigate to a directory on your file system.

For platforms other than z/OS HFS Deployment, in addition to using the selector button, you can enter the full path of the physical directory to be mapped, in the format required on your platform. (If there are spaces in the directory you are mapping, you must enclose the entire path in double quotation marks.)

For z/OS HFS and PDS servers, in addition to using the selector button, you can enter values using the following formats:

```
ext=//DD:ddname[ ; ext2=//DD:ddname2][...][ ; extn=//DD:ddnamen]
```

where *extn* are file type extensions.

```
/dir/subdir
```

entered manually or using the selector button.

For z/OS PDS servers, you can also inform the high-level qualifiers of the data set collection that comprise this application. Here is an example of the format:

```
IADMIN.SRV90.MAPAPP
```

where the user has the following datasets (not in approot):

```
IADMIN.SRV90.MAPAPP.FOCEXEC.DATA  
IADMIN.SRV90.MAPAPP.MASTER.DATA  
IADMIN.SRV90.MAPAPP.ACCESS.DATA
```

5. Select *New application (directory will be created)* from the *Map* to drop-down menu.
6. Optionally, enter a description in the *Description* field.
7. *Add directory to APPPATH* is the default. Optionally, you can decide not to add the directory.
8. Select a position from the *Position in APPPATH* drop-down menu. The options are *Last* and *First*. The default is *Last*.
9. Select a profile from the *Profile* drop-down menu.

For Server Administrators, the default is *edasprof*. For all other users, the default is the user profile.

10. Click *OK*.

The mapping is added to the *Application Directories* folder.

Procedure: How to Delete an Application or the Application Mapping

Applications and application mappings can be deleted from either the Reporting Server browser interface or the Data Migrator desktop interface.

1. From the Reporting Server browser interface menu bar, go to the *Applications* page, or from the Data Migrator desktop interface, expand the *Server* node folder.
2. Right-click the application or application mappings and select *Delete* or *Delete Mapping*, respectively.

A confirmation dialog box opens.

3. Click *OK* to delete the application or application mapping.

Note: Deleting an application mapping will delete it from any APP PATH commands that reference it, if the APP PATH command is in the same profile as APP MAP.

Indexing Applications For Global Searches

The Search Index feature is available from the WebFOCUS Reporting Server browser interface. This feature provides a single, unified search across both the WebFOCUS Client and Reporting Server repositories. This single index of results is governed by the permissions and rights for users leveraging the existing index architecture. Once your content and data are indexed, all subsequent changes to the repository content and data directories that are updated using the WebFOCUS UI or HOLD file processing are indexed automatically by WebFOCUS.

Indexing Options

You have two main options for indexing:

- Indexing all applications.
- Indexing all applications in a path.

As a filter, you can choose either of these options, in addition to applications already marked as indexed mapped directories or *ALL* mapped directories that are not marked.

In the Application Directories area:

- From the Applications node, you can choose search index options to index all or some of the items from the Reporting Server.
- From the Properties dialog box of a mapped application directory, you can include a mapped directory in the search index using a check box. You can choose this option in addition to choosing search index options.
- From an application directory node, you can create an index, update an index, or delete an index. This option is available only for a mapped application directory.

Search Index Feature

The Search Index feature is available from the WebFOCUS Reporting Server browser interface. This feature provides a single, unified search across both the WebFOCUS Client and Reporting Server repositories. This single index of results is governed by the permissions and rights for users leveraging the existing index architecture. Once your content and data are indexed, all subsequent changes to the repository content and data directories that are updated using the WebFOCUS UI or HOLD file processing are indexed automatically by WebFOCUS.

Indexing Options

You have two main options for indexing:

- Indexing all applications.
- Indexing all applications in a path.

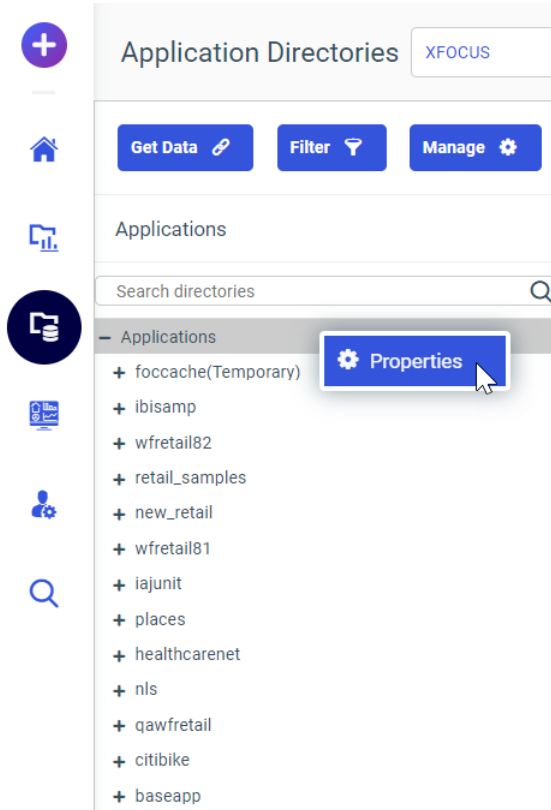
As a filter, you can choose either of these options, in addition to applications already marked as indexed mapped directories or *ALL* mapped directories that are not marked.

In the Application Directories area:

- From the Applications node, you can choose search index options to index all or some of the items from the Reporting Server.
- From the Properties dialog box of a mapped application directory, you can include a mapped directory in the search index using a check box. You can choose this option in addition to choosing search index options.
- From an application directory node, you can create an index, update an index, or delete an index. This option is available only for a mapped application directory.

Procedure: How to Choose Search Index Options for the Applications Node

1. From the Application Directories area, right-click the *Applications* node, and then select *Properties*, as shown in the following image.



The Search Index Options dialog box opens, with options for files and location, as shown in the following image.

Search Index Options

Files

- Files from all application directories
- Files only from application directories in defined application path

Location ?

- All regular directories & marked mapped directories
- All mapped directories

Cancel Save

2. Select search indexing options.

You can select the following Search Index Options for application directory files:

- Files from all application directories.** Indexes all files on the WebFOCUS Reporting Server.
- Files only from application directories in defined application path.** Indexes only the files that appear in the application directories in the defined application paths.

You can select the following Search Index Options for the application directory location:

- All regular directories & marked mapped directories.** Indexes application directories under the APPROOT and the application directories with the *Include in search index* option enabled.
- All mapped directories.** Indexes every mapped directory regardless of whether it has the *Include in search index* option enabled.

3. Click Save.

Your Reporting Server data is now indexed.

Procedure: How to Include a Mapped Directory in the Search Index Using the Properties Dialog Box

1. From the Application Directories area, under the Applications node, right-click an *Application Directory*, and select *Properties*.

The Properties dialog box opens with the Include in search index check box, as shown in the following image.

The screenshot shows a 'Properties' dialog box with the following fields and values:

- Application Type:** Application Mapping
- Physical Location:** \tibco.com\foc\data\ms2014_retail
- Map to:** New application (directory will be created)
- Mapping defined in profile:** EDASPROF
- APPPATH defined in profile:** EDASPROF
- Number of files:** 13
- Include in search index:** (highlighted with a red box)

Buttons at the bottom: Cancel, Save

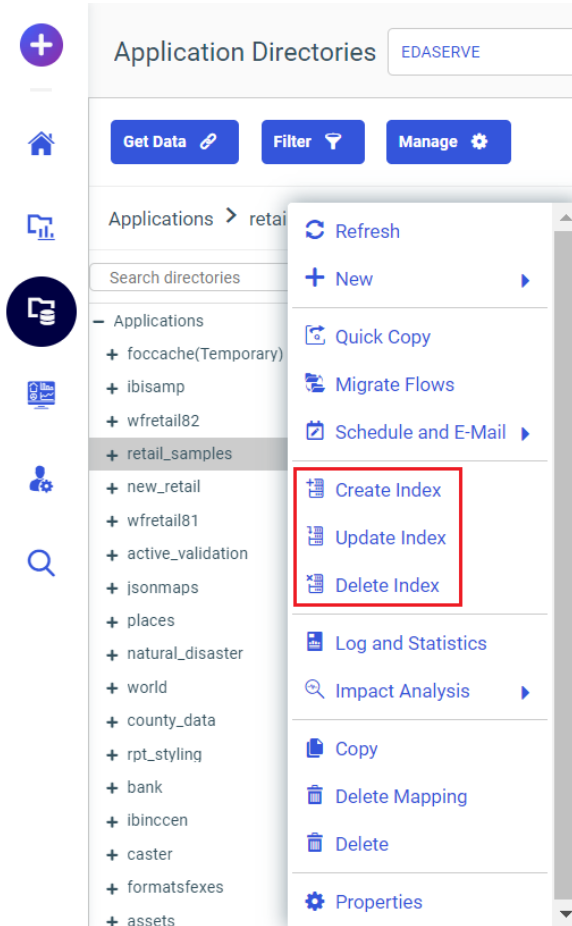
2. Select the *Include in search index* check box.

Procedure: How to Create, Update, or Delete an Index for a Mapped Application Directory

1. From the Application Directories area, under the Applications node, right-click an *Application Directory*.

The Create Index, Update Index, and Delete Index options appear on the right-click menu, as shown in the following image.

Note: These options only appear for mapped application directories.



2. Select a search index option.

The search index options are:

- Create Index.** Allows you to create an index for the selected mapped application directory.
- Update Index.** Allows you to update an existing index for the selected mapped application directory.
- Delete Index.** Allows you to delete an existing index for the selected mapped application directory.

Using an SQL Database to Store Application Contents

In addition to using a physical location, the contents of an application folder can also be stored in an SQL database. To use an SQL database, you must first create a new SQL Repository or link to your WebFOCUS Client repository. After the repository is created, you can create applications mapped to the repository and store files there.

Warning: It is advisable to create a separate SQL Connection to use in creating an SQL Repository. If an existing SQL Connection is used, only Server Administrators or users with WSCFG privileges will be able to use the synonyms from this SQL connection in the application. This provides protection for the SQL Repository, preventing unauthorized users (without administrator privileges) from accessing the contents of the SQL Repository through a synonym.

Procedure: How to Create an SQL Repository to Store Applications

You must have an adapter connection configured to an SQL database.

1. From the Applications page, click *Manage*, then *External Repository*, then *SQL Repository*, and then select *Create New*.

The Create New SQL Repository page opens, as shown in the following image.

Adapter: MS SQL Server ODBC

Connection: LocalSQL

Prefix: Id

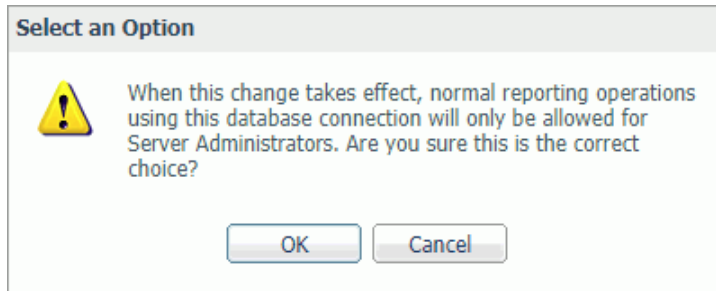
Overwrite existing repository tables and their synonyms

This will create SQL repository catalog tables on this connection and the 2 synonyms describing the catalog tables in EDACONF/catalog/IOH .

OK Cancel

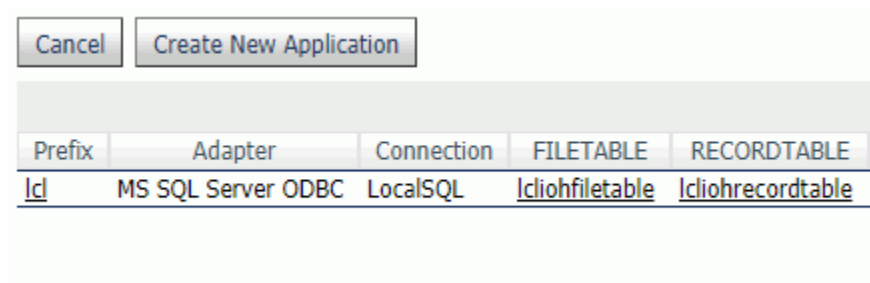
2. Select a configured adapter from the Adapter drop-down menu.
3. Select a connection from the Connection drop-down menu.
4. Enter a prefix in the *Prefix* field.
5. Optionally, select the *Overwrite existing repository tables and their synonyms* check box.
6. Click *OK*.

A warning message is displayed, as shown in the following image.



7. Click *OK*.

Two SQL Repository catalog tables are created with this connection, as shown in the following image.



The tables are:

`prefixIOHFILETABLE`

`prefixIOHRECORDTABLE`

Two synonyms describing the catalog tables are also created in EDACONF/catalog/IOH.

8. Optionally, click *Create New Application*.

The Create New Application page opens.

***Procedure:* How to Create an Application With SQL Content**

You must have created an SQL Repository.

1. From the Applications page, click *New*, then *Application Directory*.

The Create New Application page opens.

2. Select a New Application page with the Application Type set for the repository for your DBMS and connection, as shown in the following image.

The screenshot shows a 'Create New Application' dialog box. The title bar includes '(CON1)' and 'Create New Application'. The dialog contains the following fields and options:

- Application Type:** Oracle (CON1)
- Application Name:** sql01
- Physical location:** SQL:/repos
- Map to:** Existing application
- Description:** (empty text area)
- Add directory to APPPATH
- Position in APPPATH:** Last
- Profile:** edasprof
- Application IO Attributes:** (dropdown menu)

At the bottom right, there are 'Cancel' and 'OK' buttons.

Note: The choices will include the repository, adapter type, and connection name. In this example, they are the repos repository, Adapter for Oracle, and CON01 connection.

3. Enter or select values for the following parameters:
 - Application Name.** Accept the default or type a new name.
 - Physical location.** You can accept the value automatically configured or click the ellipsis (...) to navigate to the location.

- Map to.** Select Existing application, New application (directory will be created), or Existing application re-created (all files will be deleted).
 - Description.** Optionally, enter a description in the corresponding field.
 - Add directory to APPPATH.** Optionally, deselect the *Add directory to APPPATH* check box. The application is added to the APPPATH by default.
 - Position in APPPATH.** Select a position for the application from the *Position in APPPATH* drop-down menu. The choices are *Last* or *First*. The default value is *Last*.
 - Profile.** Select a profile from the *Profile* drop-down menu. For server administrators, the default value is *edaprof*. For non-administrators, their user profile is the default value.
4. Click *OK*.

The application is added to the navigation tree. You can now use this application to store procedures, synonyms, data files, and other content.

Note: Repository procedures can be executed, and Master and Access Files stored in a repository can be accessed, during server, group, and user profile execution after all DBMS connections and an APP MAP command to the repository have been executed.

Linking to Your TIBCO WebFOCUS Client Repository

You can map an application directory to access your TIBCO WebFOCUS Client Repository.

In order to map an application to your WebFOCUS Client Repository, you must have the WebFOCUS Client Adapter for REST configured, as described in the *TIBCO WebFOCUS® Adapter Administration* manual.

Procedure: How to Link to an Existing TIBCO WebFOCUS Client Repository

1. From the Applications page, click *Manage* on the ribbon, then *External Repository*, then *WebFOCUS Client Repository*, then *Link to Existing*.

You can also right-click the Application Directories tree and click *Manage*, then *External Repository*, then *WebFOCUS Client Repository*, then *Link to Existing*.

The WebFOCUS Client Repository connection page opens, as shown in the following image.

Connection:

Prefix:

This will configure this server to use existing connection to WFC repository as a regular file device.

2. Enter a prefix or accept the default value.

The prefix can be any string but, by default, it is set to the machine name running the WebFOCUS Client. This name is derived from the WebFOCUS Client REST Adapter connection.

3. Click *OK*.

The New Application pane opens with the Application Type set to map to your WebFOCUS Client REST connection, as shown in the following image.

Application Type: Application Mapping to wfcomputer - WebFOCUS Client REST (CON01)

Application Name: wfc01

Physical location: WFC:/wfcomputer

Map to: Existing application

Add directory to APPPATH

Position in APPPATH: Last

Profile: edasprof

OK Cancel

4. Enter or select values for the following parameters.

Application Type

Select *Application mapping to hostname - WebFOCUS Client REST (CON01)*

Application Name

Enter an application name for the mapped repository application, or accept the default name.

Physical location

The default WebFOCUS Repository location will be automatically entered based on your entries on the first connection screen. However, you can click the ellipsis to browse for a location.

Map to

Select one of the following mapping selections.

- Existing application.** The Repository directories will be added to an existing application directory.
- New application (directory will be created).** A new directory for the WebFOCUS Client Repository files will be created.

- Existing application, recreated (all files will be deleted).** An existing application directory will be deleted and recreated with the WebFOCUS Repository files.

Add directory to APPPATH

This check box is selected, by default, so the application will show on the Application Directories tree.

Position in APPPATH

By default, the position is Last. You can select First.

Profile

Select a profile from the drop-down list. The default is edasprof.

5. Click *OK*.

The application is added to the Application Directories tree. You can manage WebFOCUS Client Repository files using this application directory.

Procedure: How to Delete a Link to a TIBCO WebFOCUS Client Repository

1. From the Applications page, click *Manage* on the ribbon, then *External Repository*, then *WebFOCUS Client Repository*, then *List and Delete*.

You can also right-click the Application Directories tree and click *Manage*, then *External Repository*, then *WebFOCUS Client Repository*, then *List and Delete*.

A pane opens listing the TIBCO WebFOCUS Client Repositories that are linked as applications.

2. Right-click a host name in the Prefix column, or click the down arrow in the Prefix column, and click *Delete reference to remote repository*.

A dialog box opens asking you to confirm that you want to delete the reference.

3. Click *OK*.


The application directory will not be deleted, but the reference to the TIBCO WebFOCUS Repository will no longer be available to the server.

Linking to a Google Drive Repository

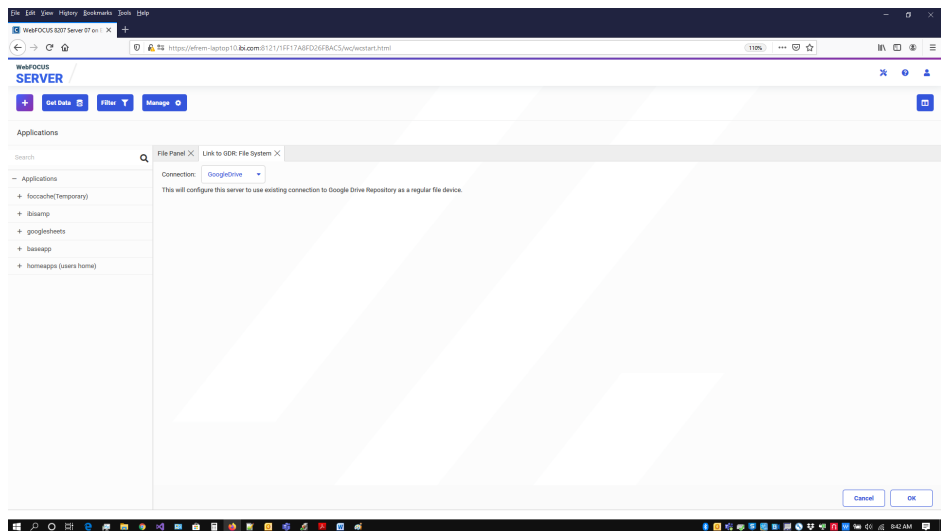
If you have configured the Adapter for Google Drive, you can link to a Google Drive Repository.

Procedure: How to Link to a Google Drive Repository

1. From the Application view of the Reporting Server browser interface, click *Manage*, then *External Repository*, then *Google Drive Repository*, then *Link to Existing*, as shown in the following image.

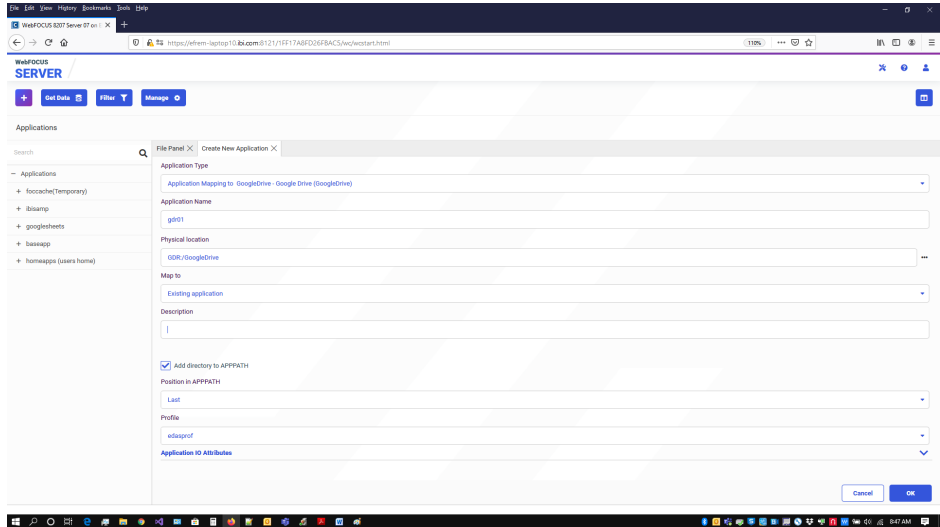


2. Choose the connection name for Google Drive, as shown in the following image.



3. Click *OK*.

A Google Drive application configuration screen opens, as shown in the following image.



4. Perform the following steps:
 - a. Enter a new Application Name and, optionally, a Description for the application.
For example:
 - Application Name:** googledrive
 - Description:** Google Drive
 - b. From the *Map to* drop-down box, select *New application (directory will be created)*.
5. Click *OK*.

Depending on the Internet speed and the number of files on the Google Drive, the mapping process could take a number of minutes to complete.

When complete, the message *Command executed successfully* will appear.

Nested Application Directories

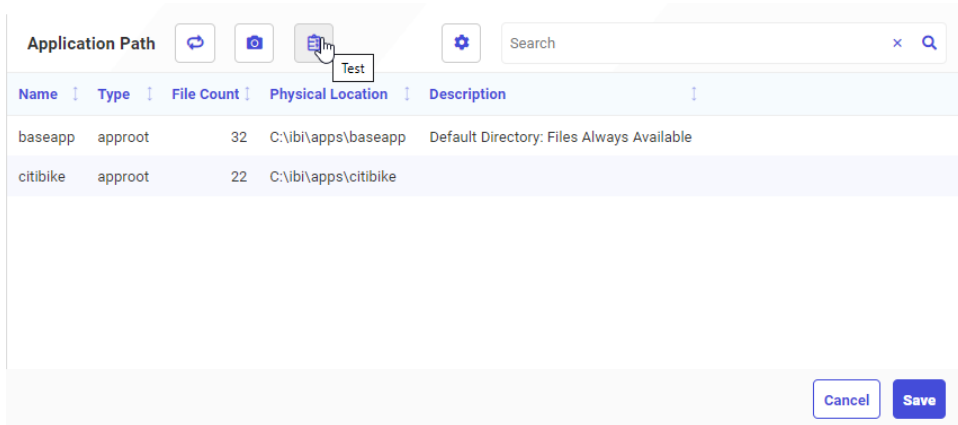
A nested application directory is one created within a higher-level application. The server allows five levels of nested application directories by default. The server must be configured for deeper or unlimited levels.

Nested application directories are implicitly added to the application path if the parent directory is on the application path.

The APP PATH command explicitly places `ibisamp` on the application path:

```
APP PATH baseapp ibisamp
```

However, you can test the path to see all of the implicitly added directories by clicking *Manage*, then *Application Path*. When the Application Path page opens, click *Test* in the Application Path pane, as shown in the following image.



The effective application path includes nested applications, such as *ibisamp/dimensions*, as shown in the following image.



Procedure: How to Set the Level of Nested Application Directories

1. From the Applications page ribbon, click *Manage*, then *Settings*.
The Application Settings pane opens.

2. Enter the level of nested applications in the *nested_app* field or select *y* from the drop-down menu. The default value is 5. Selecting *y* allows unlimited levels of nested applications.

The screenshot shows a configuration window for Home Applications. It contains several input fields and a dropdown menu:

- foccache_dir**: An empty text input field with a search icon.
- foccache_maxage**: A text input field containing the value "180".
- foccache_maxsize**: A text input field containing the value "0".
- nested_app**: A dropdown menu currently showing "5". Below it is the instruction "Select 'y' or 'n' or enter a number".
- upload_allowed**: A text input field containing a list of file extensions: "csv,ftm,dat,tab,tsv,tmp,xls,xlsx,xlsm,xlt,xlsx,xltx,xlsm,xlt,xlsx,xsd,ofd,json,c".
- upload_maxsize**: A text input field containing the value "0".
- homeapps**: A text input field containing the path "c:\1091\ib\homeapps".

At the bottom of the window are three buttons: "Restore Defaults", "Cancel", and "Save".

Note: For z/OS servers, this setting is only applicable to directory-style applications. It is not applicable to PDS-style applications or to applications mapped as a collection of dnames.

3. Click Save.

After the server restarts, you can create a new application subdirectory by right-clicking an application folder and selecting *New* and then *Application Directory* from the context menu.

Note: Nested applications must be in effect in order to create user home application directories. Some other server features also require nested application directories. For example, when you Upload a data file, the file is initially staged in a temporary upload directory under foccache while you make any edits or enhancements needed to the synonym. Tutorials use nested applications as well.

Procedure: How to Create a Nested Application Directory

1. On the Applications page, right-click an application, click *New*, and then *Application Directory*.

The Create New Application Directory page opens.

2. Enter a name in the *Application Name* field.
3. Optionally, enter a description in the *Description* field.
4. Click *OK*.

The nested application is added to the application tree under its parent application folder.

Home Application Directories for Users

The server administrator can configure the server to allow each user to have a private application directory. This directory is sometimes referred to as *myhome*, but is actually a directory under the *homeapps* directory that is named for the connected user ID (*~userid*). Providing an application directory for each user ID gives each user a directory where the user has full control to create, change, and run his applications.

The user application for any user who is not a server administrator appears in two places on the Reporting Server browser interface application tree. Both of these applications point to the same physical location. The two applications are:

- Application *~userid* (of the connected user), which is prepended to *APPPATH*.
- Application *homeapps*, which is only available to users registered under the server administrator role, is appended to *APPPATH* and can be expanded to see the user home application with the user ID name.

Server administrators can see the home directories of all users. Other users can see only their own home directories (labeled as *myhome*), unless other users share their home directories with them.

The files created in home applications can be referenced in procedures in the following ways:

- If you want to use the name *myhome* as an application name, map the name *myhome* to the internal variable named *&APPMYHOME*:

```
APP MAP myhome &APPMYHOME
```

Then, you can use *myhome* as an app reference, as in the following:

```
EX myhome/fexname
```

You can also use the variable *&APPMYHOME* as the app reference.

To make this mapping available for all users, you can place it in the edasprof server profile.

Note that the extension (.fex) is optional when executing a procedure.

- ❑ To run a request in the home directory of a specific user, you need to reference the specific user ID. Also, when a server user is registered, the primary security provider name is prepended to the user ID, and this is reflected in the app reference for a specific user ID.

The syntax includes the security provider name and the user ID, separated by an underscore. For PTH security and user1, issue the following command:

```
EX ~pth_user1/fexname
```

For more information about security providers and user IDs, see [TIBCO WebFOCUS Reporting Server Security](#) on page 35.

- ❑ You can use the procedure name with no app reference, as the home directories are automatically on the search path.

```
EX fexname
```

Note that if a procedure with the same name exists in multiple applications, you may not be running the one you want.

The first type of reference can be ported easily to any user. It enables you to create a common application that utilizes data stored in the home applications of users as myhome/data, so that each user can run the same procedure but get a report based on the data stored in the home application of that user.

The second type of reference enables you to run applications specific to a user. This type of reference can be used for testing applications before moving them to common application folders.

Home application directories should be enabled only on secured servers. If the server runs with security OFF, all users have total control of files in all applications, and the home directories will not work as designed. Nested applications must be enabled in order to create user home application directories.

The Server Administrator can monitor and manage home application directories for all users and, therefore, all user home directories are visible and in the path on the Reporting Server browser interface application tree when the connected user has Server Administrator privileges.

When the server is enabled for home applications, user home application directories are not created automatically. Users can create them from the Reporting Server browser interface Application page, or ask their server administrator to create them. The ~userid application will be automatically created for a user on the first attempt to write to it, for example, if the user uploads files to the application.

Home applications can also be stored in an SQL Repository. You must first create an SQL Repository and configure the application settings *homeapps* parameter to point to the SQL Repository, as described in [Using an SQL Database to Store Application Contents](#) on page 143.

After a home application directory is created for a user:

- ❑ The home application is added to the users application path by the server implicitly. A user cannot remove it from the path using APP commands.
- ❑ The home application directory is always first in the users active application path.
- ❑ Non-server administrator users will only see their own home application on the Application tree. The server administrator or a user with server administrator privileges will see all user home applications. The server administrator has full privileges to manage user home applications.
- ❑ The home application directory for each user will contain a file named history.txt that contains a list of recent files opened by the user. By default, the history file contains the 10 most recent files opened by the user. You can change this value. To change the number of recent files, or to disable saving recent files, on the Workspace page under Settings, click *FOCUS Sets and Info*, then *Core Engine Settings*. Under APP Service, the setting HISTORY_DEPTH configures the number of files to be kept on the recent files list. The value zero (0) means no history will be kept. You can open the list of recent files by clicking *Filter* on the Application page ribbon and selecting *Recent*.

Procedure: How to Manage Home Application Directories

In order to set up home directories for individual users, nested applications must be enabled. They are enabled to five levels by default. The home directory under which user application directories will be created is set during installation. The default is *homeapps*. The *homeapps* directory can be changed to a different physical location on the Application Setting page. All user home directories will be nested under the home directory.

1. From the Applications page, click *Manage*, then *Settings*.

The Application Settings pane opens.

2. Check the location of the *homeapps* directory. You can set the *homeapps* parameter to point to a physical directory or to an existing SQL Repository.

3. Click *Save and Restart Server* to implement these changes.

When the server restarts, go to the *Applications* menu. For a user with Server Administrator privileges, a *users home* folder will appear on the resources tree and will have all of the user home directories.

Users who do not have server administrator privileges can see their home directory under application *myhome*, which is prepended to APPPATH.


Server administrator users have the *myhome* application prepended and the *homeapps* application appended to APPPATH, The *homeapps* folder can be expanded to show the home applications for all users.

Sharing Files and Applications

Users can share their personal content under their *myhome* application with other users, groups, or roles.

To share a resource under your *myhome* application:

1. Right-click the resource you want to share, and click *Share* on the context menu. The *Manage Share Privileges* page opens, as shown in the following image.



2. Click the plus sign (+) and select *User*, *Group*, or *Role* from the *Share for* drop-down list, as shown in the following image.

The image shows a 'Share' dialog box with the following elements:

- Title: **Share** (with a close button 'X')
- Path: myhome/app01/
- Label: Share for (with a help icon '?')
- Dropdown menu: Opened, showing options: Group, User, Role, View. 'View' is selected.
- Text: access permissions
- Buttons: Cancel, Next

3. Select the security provider from the *Security Provider* drop-down list.
4. Select the type of access allowed, either *View* or *Edit*, from the *Access* drop-down list.
5. Click *Next*.

The *Share* dialog box opens, as shown in the following image.

6. Check the boxes for the users, groups, or roles that will share this resource, as shown in the following image, and click Save.

The screenshot shows a 'Share' dialog box with the following details:

- Share path: myhome/retail1.fex
- Share for: ROLES
- Search bar: Search
- Roles list:
 - Role
 - Server Administrator
 - Application Administrator
 - Server Operator
 - Basic User
 - None
 - custom1
- Buttons: Cancel, Save

7. To manage the shared resource, right-click it and click *Share*. The *Manage Share Privileges* page opens, showing the share properties that have been configured.

You can edit or delete existing share properties, add additional share properties, save your updates, or reset the page.

When a user with share privileges signs into the server, an application named *shared with me* (*users home*) is added to the application tree and provides access to the shared resources.

You can manage shared resources by clicking *Manage*, then *Shared Resources* on the Applications page.

Procedure: How to Store Home Applications in an SQL Repository

You must first create an SQL Repository. For information, see [How to Create an SQL Repository to Store Applications](#) on page 143.

1. From the ribbon, click *Manage*, then *Settings*.

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

The screenshot shows the Application Settings pane with the following fields and values:

- focache_dir**: [Empty text box]
- focache_maxage**: 180
- focache_maxsize**: 0
- nested_app**: 5 (with a dropdown arrow)
- upload_allowed**: csv,ftm,dat,tab,tsv,tmp,xls,xlsx,xltm,xltbx,xlsm,xlt,xml,xsd,dt,json,c
- upload_maxsize**: 0
- homeapps**: c:\1091\ib\homeapps

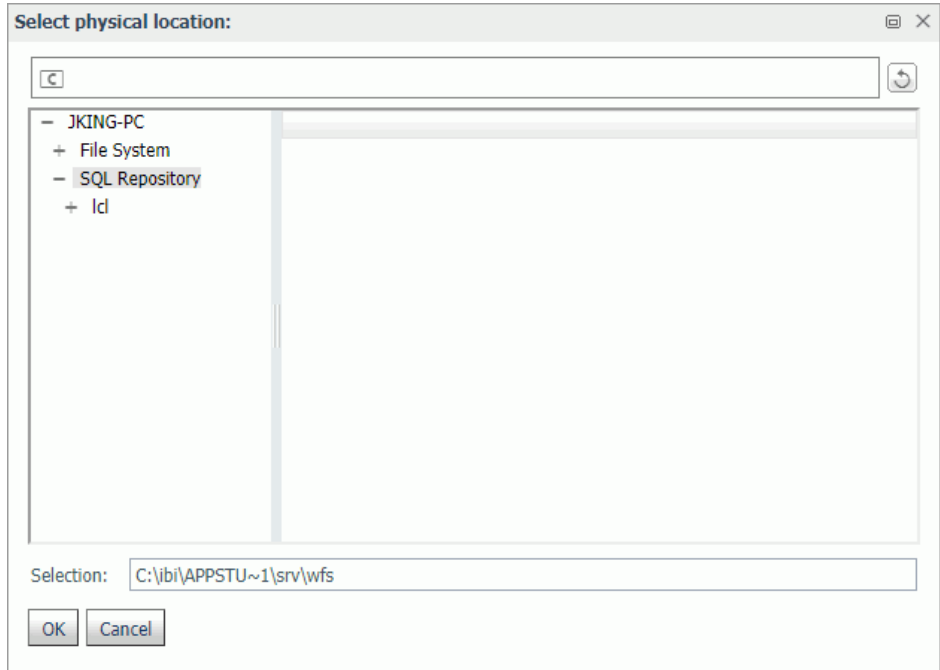
Below the fields, there is a text block: "Home Applications feature allows each connected user to keep private versions of hold files, synonyms and other WebFocus objects. Location of home application root can be on Disk (type or use File picker to set physical directory) or SQL (type SQL:/repository_prefix). Note: nested_app must not be 'n' to enable this setting."

At the bottom, there are three buttons: "Restore Defaults", "Cancel", and "Save".

2. Click the selector button (...) next to the homeapps field.

The Select physical location dialog box opens.

3. Select *SQL Repository*, as shown in the following image.



4. Select a subfolder, and click *OK*.

The subfolder is entered in the *homeapps* field, as shown in the following image.



5. Click *Save*.

Home applications that are created will now be stored in the SQL database.

Configuring the Application Path

The applications available for inclusion in the search path are identified by name, type, and physical location. If the *Add directory to APPPATH* check box was selected when the application was created, it was automatically added to the search path. If not, you must *explicitly* add it to your search path.

Note: You can also create profiles from the Application Path page. Profiles are the locations in which the search path is saved.

Procedure: How to Configure the Application Path

You can configure the Application Path to add or remove applications or mappings from either the Reporting Server browser interface or the Data Migrator desktop interface.

The User Interface for configuring the application path uses a double list.

You configure the application path from the Applications page on the Reporting Server browser interface.

1. Click *Manage* on the ribbon, then *Application Path*.

The Application Path Configuration page opens as a double-list table. The left pane shows a list of all available applications and the right shows all applications in the active Application Path, as shown in the following image.

Available Applications					Application Path				
Name	Type	File Count	Physical Location	Description	Name	Type	File Count	Physical Location	Description
app01	aproot	1	C:\lb\apps\app01	ASesri	libisamp	aproot	167	C:\lb\apps\libisamp	
app02	aproot	0	C:\lb\apps\app02		baseapp	aproot	11	C:\lb\apps\baseapp	Default Directory: Files Always Available

2. To add an application to the application path, drag it from the Available Applications list to any position on the Application Path list.

You can move applications up or down on the path by dragging them up or down on the list.

3. To remove an application from the application path, drag it from the Application Path list to the Available Applications list.
4. When you have finished, click *Save*.

On both lists, you can:

- Search for applications using the Search.
- Customize the columns displayed, or reset the display to its default, using the View icon.

On the Application Path menu bar, you can:

- Click *Switch Profile* to open the Application Path Configuration - Select Profile dialog box and select a different profile from the drop-down list.
- Click *Preview* to preview the commands in the selected profile
- Click the *Test* button to view the APP PATH command and the effective application path.

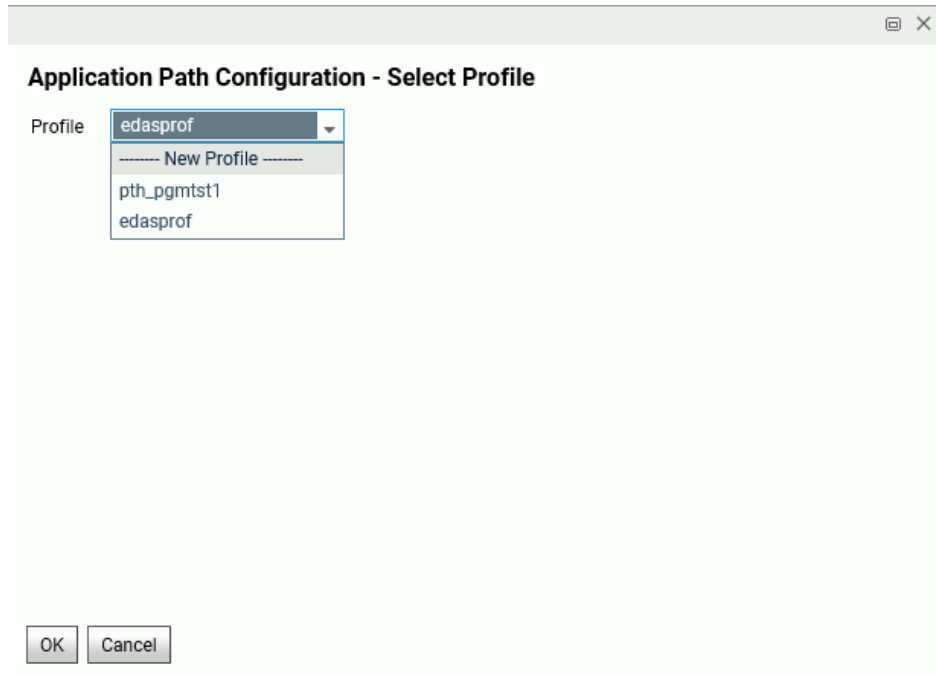
The navigation pane is updated.

Note: You can create a new profile from the Application Path Configuration - Select Profile dialog box.

Procedure: How to Configure the Application Path in a User, Group, or Role Profile

The Application Path can be configured from either the Reporting Server browser interface or the Data Migrator desktop interface.

1. From the Applications page, click *Manage*, then *Application Path*.
The Application Path page opens.
2. From the *Profile* drop-down menu, select *New Profile*.



3. Enter a name in the *New Profile Name* field and click OK.
4. Select an option from the *Change Profile Precedence* drop-down list on the menu bar.
The options are:
 - Inherit from previously executed profiles*
 - Override previously executed profiles*
 - Prepend previously executed profiles*
 - Append to previously executed profiles*
5. Optionally, click *Preview* to see the profile.

6. Click *Save*.

Procedure: How to Edit a Profile

1. Go to the *Workspace* page.
2. On the navigation pane, open the *Configuration Files* and *User/Group Profiles* folders.
3. Right-click the profile and select *Edit*.

The profile opens in a text editor with its current path displayed.

4. Edit the path information and click the *Save* icon.

Tip: You can also edit a profile search path by selecting and saving configuration options. Follow instructions for configuring the application path.

Filtering the Application Tree

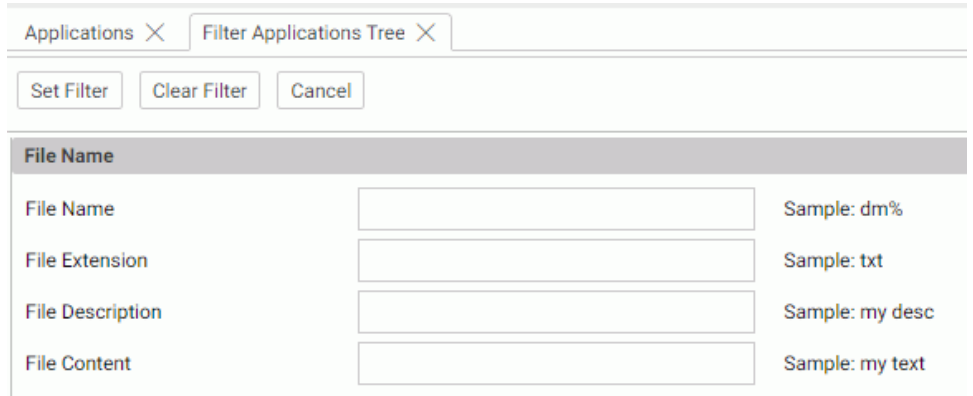
Filtering enables you to customize the file listings on the Application tree, displaying only the files you choose. It can be based on file name, location, statistics, type, or any combination of items matching a number of criteria.

Procedure: How to Filter Items on the Application Tree

You can customize the items that display on the Application Tree by filtering. Your filtering selections apply to all applications displayed in the navigation pane.

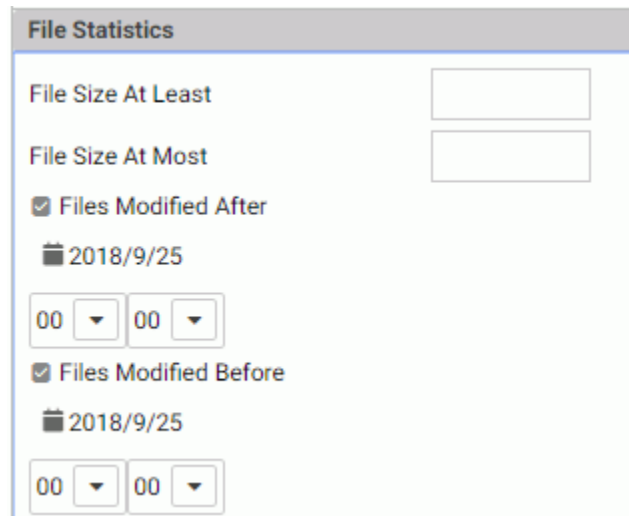
1. Go to the *Applications* page.
2. Click *Filter*, and select *Procedures*, *Synonyms*, or *Advanced*.
 - a. If you choose *Procedures*, you can display all procedures or only scheduled ones.
 - b. If you choose *Synonyms*, you can display all synonyms or only cluster or business views.
 - c. If you choose *Advanced*, the Filter Applications Tree page opens.

Using the File Name section, you can filter by name, extension, description and content, as shown in the following image. To filter the name, you can use the percent sign (%) as a wildcard character. Specifying e% displays all files whose name begins with the letter e.



File Name		
File Name	<input type="text"/>	Sample: dm%
File Extension	<input type="text"/>	Sample: txt
File Description	<input type="text"/>	Sample: my desc
File Content	<input type="text"/>	Sample: my text

Using the File Statistics section, you can filter by file size and modified date, as shown in the following image.



File Statistics	
File Size At Least	<input type="text"/>
File Size At Most	<input type="text"/>
<input checked="" type="checkbox"/> Files Modified After	<input type="text" value="2018/9/25"/>
	00 <input type="text"/>
	00 <input type="text"/>
<input checked="" type="checkbox"/> Files Modified Before	<input type="text" value="2018/9/25"/>
	00 <input type="text"/>
	00 <input type="text"/>

Using the File Type section, you can filter by type of file to be included, as shown in the following image.

File Type

Metadata

Adapters

Cluster Only

Business View Only

Procedures

Procedures

Flows

Regular Flows with IUD

DBMS SQL Flows

Direct Load Flows

Direct Load Flows with IUD

Custom Copies


Compound Reports


Shared Metadata


Cube Browser


User Functions


Scheduled Only


 (mas)


Cluster Only  (mas)

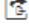
Business View Only  (mas)


Procedures  (fex)

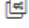
Flows  (fex)


Regular Flows with IUD  (fex)

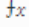
DBMS SQL Flows  (fex)


Direct Load Flows  (fex)

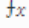
Direct Load Flows with IUD  (fex)

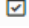
Custom Copies  (fex)

Compound Reports  (fex)

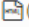
Shared Metadata  (fex)

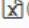
Cube Browser  (fex)

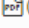
User Functions  (fex)

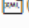
Scheduled Only  (fex)

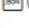
Documents


HTML  (htm, html, shtml, htt, mht, mhtml, cfm, tpl, hta, htb)

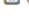
MS Excel Documents  (xls, xlsx, xlsb, xht, xltx, xlsx, xlsm, xltm)

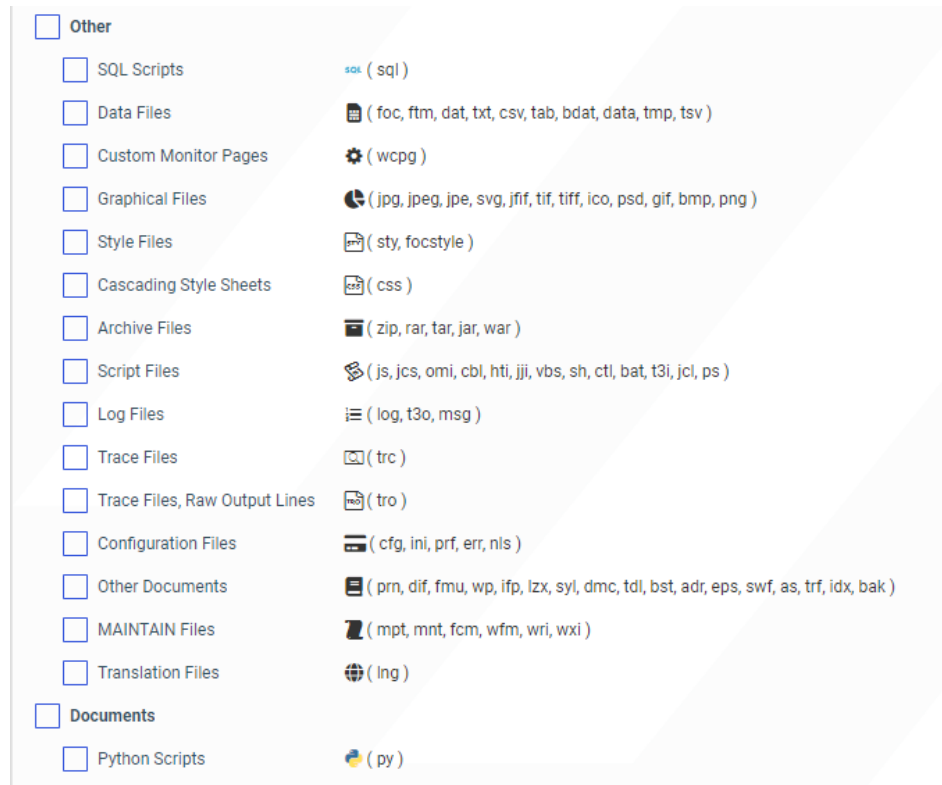
Adobe Acrobat Documents  (pdf, ai)

XML Documents  (xml, wsd, xsd, wsd, mxml, gcl, xul, dtd, xsl, xsit, axl, xhtml, xht)

JavaScript Object Notations  (json)

MS PowerPoint Documents  (ppt, pptx, pptm)

MS Word Documents  (doc, docx, docm, dot, dotx, dotm)



When you are finished, click *Set Filter*. The Filter Status page confirms that the filter was set.

You can click *Clear Filter* to clear the filter.

- d. If you choose *Impact Analysis*, you can run reports showing synonyms by procedures, procedures by synonyms, columns by procedures, or procedures by columns.
- e. If you choose *Recent*, you can see the list of files you opened recently (stored in history.txt in your myhome application).
- f. If you choose *Favorites*, you can see a list of the files you added to your favorites.

When a filter is applied, the *Application Directories* tree label includes (*Filtered*).

Note: You can remove the filter by clicking *Clear* on the Filter menu.

Searching for Files

The Search tool on the Files page provides a search function on the Application tree.

Procedure: How to Search for Files

The Search tool provides a text box for entering search criteria.

1. Go to the *Applications* page.
2. Enter a search term in the Search text box in either the navigation pane or the files pane.
3. Click the *Search* button.

The list is filtered to show the objects matching the search term, with the search term highlighted.

Sorting the Application Tree

Sorting enables you to change the order in which items are listed in an application directory.

Procedure: How to Sort an Application Directory

1. Go to the *Applications* page.
2. Open an application directory in the file panel and click a column title.

The page is sorted in descending order of the values in that column and a down arrow displays. Click again to sort in ascending order. The arrow changes to an up arrow. Click again to return to the original order. The arrow no longer displays.

Selecting Application Files

You can copy and paste files between applications, delete them from applications, and refresh synonyms in applications. You can also copy, move, or delete subfolders from one application to another. The subfolders will be copied, moved, or deleted with all of their files and all of their own subfolders.

Procedure: How to Copy, Cut, Paste, and Delete Application Files

1. Go to the *Applications* page.
2. Click an application folder to open it in the file panel.
3. On the file panel, you can select one or more files.

To select contiguous files, click the first file, hold *Shift*, and click the last file.

To select discontinuous files, hold *Ctrl*, and continue clicking the files you want to select.

4. Right-click the selected files, and click *Copy*, *Cut*, or *Delete*.
5. To paste copied or cut files into another application folder, right-click the application folder and click *Paste*.

The file(s) are pasted to the selected application.

Procedure: How to Refresh Synonyms in an Application

1. Go to the *Applications* page.
2. Right-click an application folder, point to *Metadata Management*, and click *Refresh*.

The *Update Synonyms* page opens. You can select the attributes to update. Certain attributes, such as keys, indexes, and formats are selected by default.

3. When you have made your selections, click *Next*.

A status page opens that reports the result of the update request.

Indexing Applications For Global Searches

The Search Index feature is available from the WebFOCUS Reporting Server browser interface. This feature provides a single, unified search across both the WebFOCUS Client and Reporting Server repositories. This single index of results is governed by the permissions and rights for users leveraging the existing index architecture. Once your content and data are indexed, all subsequent changes to the repository content and data directories that are updated using the WebFOCUS UI or HOLD file processing are indexed automatically by WebFOCUS.

Indexing Options

You have two main options for indexing:

- Indexing all applications.
- Indexing all applications in a path.

As a filter, you can choose either of these options, in addition to applications already marked as indexed mapped directories or *ALL* mapped directories that are not marked.

In the Application Directories area:

- From the Applications node, you can choose search index options to index all or some of the items from the Reporting Server.
- From the Properties dialog box of a mapped application directory, you can include a mapped directory in the search index using a check box. You can choose this option in addition to choosing search index options.
- From an application directory node, you can create an index, update an index, or delete an index. This option is available only for a mapped application directory.

Search Index Feature

The Search Index feature is available from the WebFOCUS Reporting Server browser interface. This feature provides a single, unified search across both the WebFOCUS Client and Reporting Server repositories. This single index of results is governed by the permissions and rights for users leveraging the existing index architecture. Once your content and data are indexed, all subsequent changes to the repository content and data directories that are updated using the WebFOCUS UI or HOLD file processing are indexed automatically by WebFOCUS.

Indexing Options

You have two main options for indexing:

- Indexing all applications.
- Indexing all applications in a path.

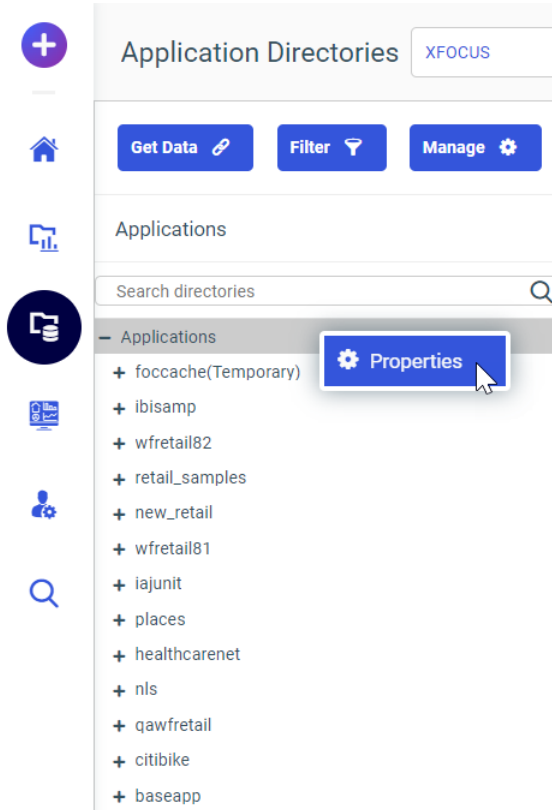
As a filter, you can choose either of these options, in addition to applications already marked as indexed mapped directories or *ALL* mapped directories that are not marked.

In the Application Directories area:

- From the Applications node, you can choose search index options to index all or some of the items from the Reporting Server.
- From the Properties dialog box of a mapped application directory, you can include a mapped directory in the search index using a check box. You can choose this option in addition to choosing search index options.
- From an application directory node, you can create an index, update an index, or delete an index. This option is available only for a mapped application directory.

Procedure: How to Choose Search Index Options for the Applications Node

1. From the Application Directories area, right-click the *Applications* node, and then select *Properties*, as shown in the following image.



The Search Index Options dialog box opens, with options for files and location, as shown in the following image.

Search Index Options

Files

Files from all application directories

Files only from application directories in defined application path

Location ?

All regular directories & marked mapped directories

All mapped directories

Cancel Save

2. Select search indexing options.

You can select the following Search Index Options for application directory files:

- Files from all application directories.** Indexes all files on the WebFOCUS Reporting Server.
- Files only from application directories in defined application path.** Indexes only the files that appear in the application directories in the defined application paths.

You can select the following Search Index Options for the application directory location:

- All regular directories & marked mapped directories.** Indexes application directories under the APPROOT and the application directories with the *Include in search index* option enabled.
- All mapped directories.** Indexes every mapped directory regardless of whether it has the *Include in search index* option enabled.

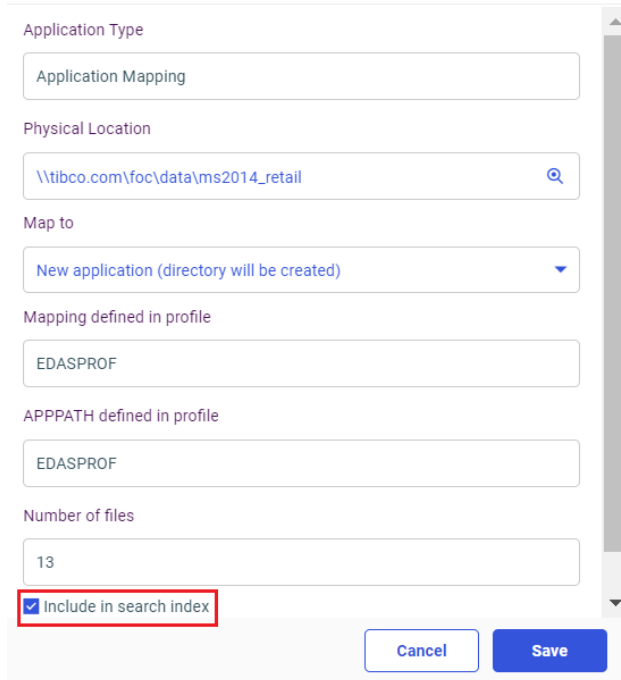
3. Click Save.

Your Reporting Server data is now indexed.

Procedure: How to Include a Mapped Directory in the Search Index Using the Properties Dialog Box

1. From the Application Directories area, under the Applications node, right-click an *Application Directory*, and select *Properties*.

The Properties dialog box opens with the Include in search index check box, as shown in the following image.



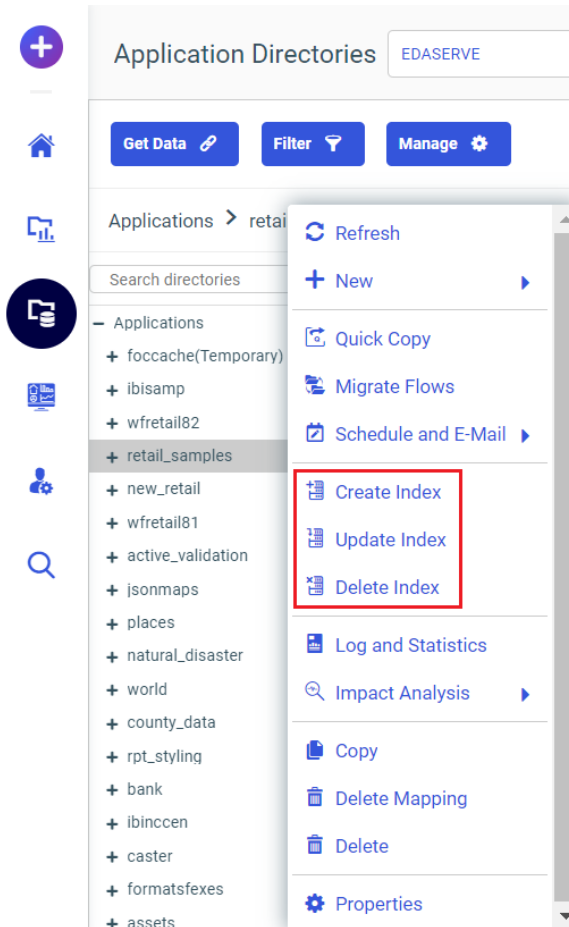
2. Select the *Include in search index* check box.

Procedure: How to Create, Update, or Delete an Index for a Mapped Application Directory

1. From the Application Directories area, under the Applications node, right-click an *Application Directory*.

The Create Index, Update Index, and Delete Index options appear on the right-click menu, as shown in the following image.

Note: These options only appear for mapped application directories.



2. Select a search index option.

The search index options are:

- Create Index.** Allows you to create an index for the selected mapped application directory.
- Update Index.** Allows you to update an existing index for the selected mapped application directory.
- Delete Index.** Allows you to delete an existing index for the selected mapped application directory.

Application Commands Overview

This topic lists the platform-independent application (APP) commands that enable you to control the application environment.

You can use the following wildcard characters in the file name and file type references for the APP commands COPYFILE, MOVEFILE, DELETEFILE, and RENAMEFILE.

- ❑ An asterisk (*) replaces any combination of characters of any length (including zero).
Note that an asterisk can also be used to replace the entire filename or filetype parameter.
- ❑ A question mark (?) replaces zero or one character.

Reference: APP Commands Quick Reference

Click any command in the following charts to access detailed information, including the required syntax.

Search Path Management Commands

Command	Description
<i>APP PATH</i>	Sets or resets the application search path.
<i>APP PREPENDPATH</i>	Temporarily adds application names to the beginning of an existing APP PATH search path.
<i>APP APPENDPATH</i>	Temporarily adds application names to the end of an existing APP PATH search path.
<i>APP MAP</i>	Defines a virtual application that points to a physical location outside of the approot structure or redirects an application to another application. This command makes the virtual or redirected application available for addition to the search path. It does not automatically add it to the APP PATH.
<i>APP SET METALLOCATION_SAME</i>	Indicates whether corresponding Master and Access files must be in the same application directory.
<i>APP ? METALLOCATION_SAME</i>	Retrieves the value of APP SET METALLOCATION_SAME.
<i>APP SHOWPATH</i>	Lists all the currently active applications in the search path.

Application Management Commands

Command	Description
<i>APP COPY</i>	Copies the contents of one application to a second application.
<i>APP CREATE</i>	Creates an application under the approot location.
<i>APP DELETE</i>	Deletes an application.
<i>APP MOVE</i>	Moves the contents of one application to a second application.
<i>APP PROPERTY CODEPAGE</i>	Specifies a code page for files in an application.
<i>APP RENAME</i>	Renames an application.

File Management Commands

Command	Description
<i>APP COPYF[ILE]</i>	*Copies a single component or component type from one application to another.
<i>APP MOVEF[ILE]</i>	*Moves a single component or component type from one application to another.
<i>APP RENAMEF[ILE]</i>	*Renames a single component or component type in an application.
<i>APP DELETEF[ILE]</i>	*Deletes a single component or component type from an application.

* The shortened form of the APP commands is used in the remainder of this document.

Output Redirection Commands

Command	Description
<i>APP HOLD</i>	Controls where output files are created for any HOLD, SAVE, SAVB, CREATE SYNONYM, or APP QUERY HOLD process in the application, unless a FILEDEF command has been used to allocate data files.
<i>APP HOLDDATA</i>	Designates an application as the location for temporary data files created with the HOLD, SAVE, or SAVB command.
<i>APP HOLDMETA</i>	Designates an application as the location for temporary Master and Access Files created with the HOLD command.
<i>APP FI[LEDEF]</i>	This command has been deprecated and aliased to FILEDEF.

Help Commands

Command	Description
<i>APP HELP</i>	Displays a list of APP commands with a brief description of each.

Reference: Application Metadata Commands and Metadata Tables

Click any command in the following chart to access detailed information, including the required syntax.

Command	Description
<i>STATE app/file.extension</i>	Checks file existence.
<i>APP LIST app [HOLD]</i>	Lists the applications under approot. If the HOLD option is used, it lists the applications under approot and writes the output to a temporary file called focappl.ftm, which you can then use in a report request.

Command	Description
<code>APP QUERY app [HOLD]</code>	Lists all files in the application. If the HOLD option is used, it lists all files in the application and writes the output to a temporary file called focappq.ftm, which you can then use in a report request.
<code>catalog/sysfiles</code>	Table, list of accessible app name objects on path for a given type (default MASTER).
<code>catalog/sysdirs</code>	Table, recursive list of physical files under a physical directory.
<code>catalog/sysapps</code>	Table, metadata for physical objects on path.
<code>catalog/systables</code>	Table, app name of tables (and related metadata) on path.

For information about reporting from system tables, see the *TIBCO WebFOCUS® Developing Reporting Applications* manual.

Search Path Management Commands

The server has a default search path for application and system components. You can supplement this search path by using one or more of the following APP commands:

- APP PATH
- APP PREPENDPATH
- APP APPENDPATH
- APP SET METALLOCATION_SAME
- APP ? METALLOCATION_SAME

Generally, these commands add applications to the beginning of the default search path. The exception is temporary components that are created in the current session. These temporary components are searched first, before the user-defined path.

You can issue the APP PATH command manually or set the application search path from the Reporting Server browser interface or the Data Migrator desktop interface. When you configure the application path from the Reporting Server browser interface or the Data Migrator desktop interface, the APP PATH command is stored in a selectable profile (global, group, or user). The global profile, edasprof, is the default.

APP PATH

The APP PATH command sets the search path to a designated list of application names that refer to applications under the approot value. You can specify multiple application names to extend the search path.

Syntax: How to Add an Application to the Search Path Manually

```
APP PATH app1[/] [app2[/] ...]  
        [appn[/]]
```

where:

```
app1 . . . appn
```

Are application names. If you follow an application name with a slash (/), nested applications (the subtree of applications below the named application) will not be in the search path. If you do not follow the application name with a slash, the *nested_app* parameter in the edaserve.cfg file determines whether nested applications are searched for files referenced in a procedure, and to what level. If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and code more application names on the next line.

Note:

- You can use the APP PATH command without an application name to reset the search path to the initial list.
- APP PATH does not validate the application list.

APP PREPENDPATH

The APP PREPENDPATH command enables you to temporarily add application names to the beginning of an existing APP PATH search path.

If you wish to use this command to alter the search path, you must code it manually in your application.

Syntax: **How to Add Application Names to the Beginning of a Search Path**

```
APP PREPENDPATH app1[/] [app2[/]] ...
                [appn[/]]
```

where:

app1 . . . *appn*

Are application names. If you follow an application name with a slash (/), nested applications (the subtree of applications below the named application) will not be in the search path. If you do not follow the application name with a slash, the *nested_app* parameter determines whether nested applications are searched for files referenced in a procedure, and to what level. If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and code more application names on the next line.

APP APPENDPATH

The APP APPENDPATH command enables you to temporarily add application names to the end of an existing APP PATH search path.

If you wish to use this command to alter the search path, you must code it manually in your application.

Syntax: **How to Add Application Names to the End of a Search Path**

```
APP APPENDPATH app1[/] [app2[/]] ... [appn[/]]
```

where:

app1 . . . *appn*

Are application names. If you follow an application name with a slash (/), nested applications (the subtree of applications below the named application) will not be in the search path. If you do not follow the application name with a slash, the *nested_app* parameter determines whether nested applications are searched for files referenced in a procedure, and to what level. If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and code more application names on the next line.

APP MAP

The APP MAP command allows you to assign an application name to a non-aproot application anywhere in the file system or to redirect an application. The application name becomes a virtual application under aproot, which can be referenced in an APP PATH command and any other APP command that takes an application name as a parameter.

Note that mapping does *not* automatically add a directory to the path, it simply makes it available for addition to the search path.

Syntax: How to Map a Physical File Location or Redirect an Application

To map a physical file location outside of aproot, the syntax is:

```
APP MAP virtualname real_location
```

where:

virtualname

Is an application name of up to 64 characters that can later be used in an APP PATH command.

real_location

Is a real full path name or DDNAME in the native style of the given operating system. On UNIX, Linux, i5 and z/OS HFS Deployment, the location may be a mixed case name, but the MAP virtualname itself is always handle insensitive when used (that is, EX mymap/mytest).

Note that if the real location contains spaces, it must be surrounded by double quotation marks.

Note: On IBM i, the APP MAP command can only be used to map an IFS directory and not a QSYS library.

To redirect an aproot or non-aproot application to a different name, the syntax is:

```
APP MAP app1 app2/dir1/dir2/./dirn
```

where:

app1

Can be an existing physical, mapped or linked app. It can also be a new app.

app2

Can be a physical, mapped or linked or non-existent app.

dir1 ... dirn

Are application directory names.

Note: Remapping of internal applications (for example, myhome, homeapps, _edaconf, _edahome) is not supported.

Example: Sample APP MAP Commands

Basic example for Windows:

```
APP MAP test c:\temptest\
```

Note that if a path name contains spaces, the name must be surrounded by double quotation marks. For example:

```
APP MAP test "c:\temp test\"
```

To remap an application under approot:

```
APP MAP ibisamp retail
```

Syntax: How to Map DDNAME Allocations

For the Unified Server (HFS and PDS deployments), the syntax for this type of mapping is

```
APP MAP appname file_extension=//dd:ddname;file_extension//  
dd:ddname;
```

where:

appname

Is the name of the application used to reference this mapping in an APP PATH, APP APPENDPATH, or APP PREPENDPATH command.

file_extension

Is one of the following valid server file extensions:

```
.mas  
.fex  
.acx  
.htm  
.sty  
.gif  
.psb
```

ddname

Is the ddname of the allocation you wish to map. The allocation can be performed using JCL code or a DYNAM command.

Example: Mapping DDNAME Allocations

```
DYNAM ALLOC FILE MYMAS DA EDAARH.MASTER.DATA SHR REU
APP MAP APP1 MAS=//DD:MYMAS;
APP APPENDPATH APP1
```

By default, the server has an APP MAP command in the edasprof.prf file to map the application MVSAPP to the allocations FOCEXEC, MASTER, ACCESS, HTML, FOCSTYLE, GIF, FOCPSB. While allocations of these ddnames are not required for the APP MAP command to be valid, once the ddnames are allocated by JCL or DYNAM commands, they become available for use.

Reference: APP MAP With Universal Naming Convention (UNC)

On platforms that support Universal Naming Convention (UNC), you must use the UNC to designate a network drive to access APP directories. The UNC must:

- Be at least one folder below the initial shared location.
- Not contain spaces unless enclosed in double quotation marks. For example,

```
\\mynode\myshare\accting
"\\mynode\my share\accting"
```

APP SET METALLOCATION_SAME

The APP SET METALLOCATION_SAME command identifies whether Master Files and their corresponding Access Files must be in the same location.

Syntax: How to Control the Location of Synonym Files

```
APP SET METALLOCATION_SAME {ON|OFF}
```

where:

ON

Specifies that Master Files and their corresponding Access Files must reside in the same application directory. ON is the default value.

OFF

Specifies that once the Master File for a request is located, the server will use the active search path to find the corresponding Access File.

APP ? METALLOCATION_SAME

The APP ? METALLOCATION_SAME command queries whether Master Files and their corresponding Access Files must be in the same location.

Syntax: How to Query Whether Synonym Files Must Reside in the Same Location

APP ? METALLOCATION_SAME

If the result of this query command is ON, the server expects to find corresponding Master and Access Files in the same application directory. If the result is OFF, the server uses the active search path to find the Access File that corresponds to a given Master File.

APP SHOWPATH

The APP SHOWPATH command lists all the currently active applications in the search path, including baseapp, which is always last. This list mirrors the list of applications displayed in the applications tree on the navigation pane.

Syntax: How to List Active Applications

APP SHOWPATH

Example: Listing Active Applications in the Search Path

The server is generally installed with two default applications: ibisamp (contains sample files), and baseapp (which can contain any files you create).

The APP SHOWPATH command generates the following output is:

```
ibisamp
baseapp
```

Application and File Management Commands

The APP commands in this section provide management options for applications and their component files.

APP CREATE

In general, the APP CREATE command creates an application under the approot location. The exception is a PDS deployment on a Unified Server, where an application is a physical entity and each of its component file types is stored in a separate PDS.

The APP CREATE command can create any number of applications with one command.

Syntax: How to Create an Application Manually

```
APP CREATE app1[/app1a...] [app2[/app2a...] ...  
          [appn[/appna...] ] [DROP]
```

where:

app1...*appn*

Are application names under approot. The application name can be up to 64 characters.

app1a...*appna*

Are nested application directories, allowed when nested applications are configured. In order to create a nested application, the parent application must already exist.

DROP

Deletes an application if one already exists with the same name as the one to be created, and then creates a new application with that name. Note that any files in the pre-existing application are deleted. Without the DROP option, a message will be generated, and the pre-existing application will not be deleted or changed.

The application name may not contain spaces. If the name contains spaces, each section is understood to be a separate application. If you require a name with spaces, you must create it using another mechanism, such as Windows Explorer. You can then use the APP MAP command to add it to APPROOT.

If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and code more application names on the next line.

The word HOLD cannot be used as an application name.

Syntax: How to Change Default Characteristics of Component File Types (PDS Deployment Only)

If you are working on a Unified Server in PDS deployment, you can change the default characteristics of individual component file types by issuing a DYNAM SET APP command. This command controls the types of component files that are generated for the application when an APP CREATE command is issued. By default, all component file types are generated.

The syntax is

```
DYNAM SET APP FOR filetype [SKIP|CREATE] [POSTFIX aaa.bbb] [parms]
```

where:

filetype

Are the component types that may be affected by this command, in uppercase: FOCEXEC, MASTER, ACCESS, HTML, GIF, FOCSTYLE, MAINTAIN, ETG. You must issue a separate command for each component type you wish to affect.

SKIP

Indicates that the designated file type should not be created when the APP CREATE command is issued.

CREATE

Creates the designated file type when the APP CREATE command is issued. This is the default setting.

POSTFIX

Specifies the lower level qualifier of the DSN (data set name) for the component type. The APPROOT value is used to complete the full DSN, which is expressed as

aprootvalue.appname.component_type

The default value for *component_type* is *filetype.DATA*.

parms

Are the allocation parameters you can set. The default parameter values are:

File Type	Parameter
FOCEXEC	RECFM VB TRKS LRECL 4096 BLKSIZE 27998 SPACE 50 50 DIR 50
MASTER	RECFM FB TRKS LRECL 80 BLKSIZE 22000 SPACE 50 50 DIR 50
ACCESS	RECFM FB TRKS LRECL 80 BLKSIZE 22000 SPACE 50 50 DIR 50
HTML	RECFM VB TRKS LRECL 4096 BLKSIZE 27998 SPACE 50 50 DIR 50

File Type	Parameter
GIF	RECFM VB TRKS LRECL 4096 BLKSIZE 27998 SPACE 50 50 DIR 50 The GIF file type creates libraries for GIF and JPG files.
FOCSTYLE	RECFM FB TRKS LRECL 1024 BLKSIZE 27648 SPACE 50 50 DIR 50
MAINTAIN	RECFM VB TRKS LRECL 4096 BLKSIZE 27998 SPACE 50 50 DIR 50
ETG	RECFM FB TRKS LRECL 80 BLKSIZE 22000 SPACE 50 50 DIR 50

Example: Changing Default Characteristics of an Application (PDS Deployment)

The following command indicates that GIF files should not be created when the APP CREATE command is issued.

```
DYNAM SET APP FOR GIF SKIP
```

The following command indicates that Procedures (FOCEXECs) should be created when APP CREATE is issued.

```
DYNAM SET APP FOR FOCEXEC TRKS SP 10 20 DIR 30
```

APP COPY

The APP COPY command copies the entire contents of one application to another. The target application must already exist.

Syntax: How to Copy an Application

```
APP COPY app1[/app1a...] app2[/app2a...]
```

where:

```
app1[/app1a...]
```

Is the application being copied. It can be a nested application name.

app2[/app2a...]

Is the application to which the contents of the first application are being copied. It can be a nested application name.

APP COPYF[ILE]

The APP COPYF[ILE] command copies one or more components or component types from one application to another.

Note that if you copy a component manually, you can, optionally, rename it in the process.

If you copy a Master File, the corresponding Access File is also copied. However, copying an Access File (file type FOCSQL) does not automatically copy the corresponding Master File.

Syntax: How to Copy an Application Component Manually

```
APP COPYF[ILE] app1[/app1a...]
  {filename1|*} filetype1 app2 [ /app2a...]
  {filename2|*} {filetype2|*} [IFEXIST] DROP
```

where:

app1[/app1a...]

Is the application that contains the component to be copied. It can be a nested application name.

filename1

Are the components to be copied. Use an asterisk (*) to copy all components of file type *filetype1*.

You can use the following wildcard characters in the file name and file type references.

- ❑ An asterisk (*) replaces any combination of characters of any length (including zero).

Note that an asterisk can also be used to replace the entire filename or filetype parameter.

- ❑ A question mark (?) replaces zero or one character.

filetype1

Is the file type, in uppercase, of the component to be copied.

app2[/app2a...]

Is the application to which the named component is being copied. It can be a nested application name.

filename2

Is the component name in the target application, after the copy process. Use an asterisk (*) to propagate the file names from the source application to the target application.

filetype2

Is the component type, in uppercase, in the target application after the copy process. Use an asterisk (*) to propagate the file types from the source application to the target application.

IFEXIST

Ignores any component in the source application that does not exist.

DROP

Overwrites any component already in the target application with the same name and file type as a component being copied.

For a full list of the types of files you can copy with APP commands, see [Designating File Types for APP Commands](#) on page 193.

APP MOVE

The APP MOVE command moves the entire contents of one application to another. The target application must already exist.

Syntax: How to Move an Application

```
APP MOVE app1[/app1a...] app2[/app2a...]
```

where:

app1[/app1a...]

Is the application being moved. It can be a nested application name.

app2[/app2a...]

Is the application to which the contents of the first application are being moved. It can be a nested application name.

APP MOVEF[ILE]

The APP MOVEF[ILE] command moves one or more components or component types from one application to another.

Note that if you move a component manually, you can, optionally, rename it in the process.

If you move a Master File, the corresponding Access File is also moved. However, moving an Access File (file type FOCUSQL) does not automatically move the corresponding Master File.

Syntax: How to Move an Application Component Manually

```
APP MOVEF[ILE] app1[/app1a...]
  {filename1|*} filetype1 app2 [/app2a...]
  {filename2|*} {filetype2|*} [IFEXIST] [DROP]
```

where:

app1[/*app1a...*]

Is the application that contains the component to be moved. It can be a nested application name.

filename1

Is the name of the component to be moved. Use an asterisk (*) to move all components of file type *filetype1*.

You can use the following wildcard characters in the file name and file type references.

- ❑ An asterisk (*) replaces any combination of characters of any length (including zero).

Note that an asterisk can also be used to replace the entire filename or filetype parameter.

- ❑ A question mark (?) replaces zero or one character.

filetype1

Is the file type, in uppercase, of the component to be moved.

app2[/*app2a...*]

Is the application to which the named component is being moved. It can be a nested application name.

filename2

Is the component name in the target application, after the move process. Use an asterisk (*) to propagate the file names from the source application to the target application.

filetype2

Is the component type, in uppercase, in the target application after the move process. Use an asterisk (*) to propagate the file types from the source application to the target application.

IFEXIST

Ignores any component in the source application that does not exist.

DROP

Overwrites any component already in the target application with the same name and file type as a component being moved.

For a full list of the types of files you can move with APP commands, see [Designating File Types for APP Commands](#) on page 193.

APP DELETE

The APP DELETE command deletes applications under approot.

Syntax: How to Delete an Application Manually

```
APP DELETE app1[/app1a...] [app2[/app2a...] ...  
          [appn[/appna...]]
```

where:

```
app1[/app1a...]... [appn[/appna...]
```

Are application names. Nested application names are supported. If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and enter additional application names on the next line.

APP DELETEDF[ILE]

The APP DELETEDF[ILE] command deletes one or more components or component types from an application.

If you delete a Master File, the corresponding Access File is also deleted. However, deleting an Access File (file type FOCSQL) does not automatically delete the corresponding Master File.

Syntax: How to Delete an Application Component Manually

```
APP DELETEDF[ILE] app[/appna...] {filename|*} filetype
```

where:

```
appn[/appa...]
```

Is the application from which the component or component type is being deleted. Nested application names are supported.

filename

Is the name of the component to be deleted. Use an asterisk (*) to delete all files of type *filetype*.

You can use the following wildcard characters in the file name and file type references.

- ❑ An asterisk (*) replaces any combination of characters of any length (including zero).

Note that an asterisk can also be used to replace the entire filename or filetype parameter.

- ❑ A question mark (?) replaces zero or one character.

filetype

Is the component type, in uppercase, of the component to be deleted.

For a full list of the types of files you can use with APP commands, see [Designating File Types for APP Commands](#) on page 193.

APP PROPERTY CODEPAGE

The APP PROPERTY appname CODEPAGE command identifies the codepage to be used for non-data files in the application directory.

Syntax: How to Specify a Code Page for an Application

```
APP PROPERTY app[/appa...] CODEPAGE number
```

where:

app[/appa...]

Is an application name. Nested application names are supported.

number

Is the code page number for non-data files in the application.

APP RENAME

The APP RENAME command renames an existing application.

Note: You cannot rename an application if it is active in the search path.

Syntax: How to Rename an Application

```
APP RENAME app1[/app1a...] app2[/app2a...]
```

where:

app1[/app1a...]

Is the application name to be renamed. It can be a nested application name.

app2[/app2a...]

Is the new application name of up to 64 characters. It can be a nested application name.

Example: Renaming an Application

The following shows app1 being renamed to app2.

```
APP RENAME app1 app2
```

APP RENAMEF[ILE]

The APP RENAMEF[ILE] command renames one or more components in an application.

If you rename a Master File, the corresponding Access File is also renamed. However, renaming an Access File (file type FOCSQL) does not automatically rename the corresponding Master File.

Syntax: How to Rename an Application Component

```
APP RENAMEF[ILE] app[/appa... ] filename1  
filename2 filetype [DROP]
```

where:

app[/appa...]

Is the name of the application that contains the component being renamed. It can be a nested application name

filename1

Is the file name of the component to be renamed.

You can use the following wildcard characters in the file name and file type references.

- ❑ An asterisk (*) replaces any combination of characters of any length (including zero).

Note that an asterisk can also be used to replace the entire filename or filetype parameter.

- ❑ A question mark (?) replaces zero or one character.

filename2

Is the new name for the component. The component name may be up to 64 characters.

filetype

Is the file type, in uppercase, of the component to be renamed.

DROP

Overwrites an existing component with the same file name and file type.

For a full list of the types of files you can use with APP commands, see [Designating File Types for APP Commands](#) on page 193.

Designating File Types for APP Commands

The APP COPYF, APP MOVEF, APP DELETEF, and APP RENAMEF commands enable you to perform their actions on a wide variety of file types.

The following is a comprehensive list of the file types you can use with APP commands and the file extensions associated with the on-disk names for hierarchical file systems.

Note that the file types must be coded in uppercase in any APP command that requires it.

Note: This list reflects file types supported across all products and release levels. Particular file types may not be supported in particular releases or with every product.

File Type	File Extension
ACX	.acx
ADR	.adr
AFM	.afm
BMP	.bmp
BST	.bst
cascading style sheet	.css
CONTROL	.ctl
DATA	.dat
DDS	.DDS

File Type	File Extension
DEFAULT	The APP <i>filename</i> value is used to derive the physical extension for the APP command, so that unknown user-defined extensions may be supported in an APP command (for example, APP COPYFILE BASEAPP MYFILE.FOO DEFAULT BASEAPP MYFILE FOCEXEC).
DTD	.dtd
EDANLS	.nls
EDAPRFU	.prf
EDAPROF	.prf
EDAPSB	.psb
EPS	.eps
ERRORS	.err
ETG	.etg
ETL	.etl
EXCEL	.xls
FMU	.fmu
FOCCOMP	.fcm
FOCDEF	.def
FOCEXEC OR FEX	.fex
FOCFTMAP	.fmp
FOCPSB	.psb
FOCSQL	.acx
FOCSTYLE	.sty

File Type	File Extension
FOCTEMP	.ftm
FOCUS	.foc
GIF	.gif
HLI	.hli
HTML	.htm
IBICPG	.sl
JPG	.jpg
JS	.js
LSN	.lsn
MAINTAIN	.mnt
MASTER OR MAS	.mas MASTER has a special behavior that any matching Access File (.acx) is also operated upon by the APP command. This is so metadata is operated upon as a matched pair. Use MAS if it is strictly desired to only operate on the Master File and not the Access File.
MHT	.mht
Microsoft Access database	.mdb
MNTPAINT	.mpt
OMI	.omi
PDF	.pdf
PFA	.pfa
PFB	.pfb

File Type	File Extension
PNG	.png
PS	.ps
SMARTLIB	.knb
SQL	.sql
SVG	.svg
TABS	.txt
TDL	.tdl
TRF	.trf
TTEDIT	.tte
TXT	.txt
WINFORMS	.wfm
WSDL	.wsd
XHT	.xht
XLSM	.xlsm
XLSX	.xlsx
XLTM	.xltm
XLTX	.xltx
XML	.xml
XSD	.xsd
XSL	.xsl

Output Redirection Commands

Three APP commands (APP HOLD, APP HOLDDATA, and APP HOLDMETA) along with the FILEDEF and DYNAM commands comprise a class of commands that control where output is stored. In order to redirect output as you wish, it is important to understand the interactions among these commands.

Note: When the same behavior applies for APP HOLD, APP HOLDDATA, and APP HOLDMETA, these commands are referred to collectively as APP HOLD*. Note also that although DYNAM (USS only) and FILEDEF are not members of the APP family of commands, these file allocation commands interact with the APP HOLD* commands. Therefore, where appropriate, these commands are also included in this discussion. APP FI[LEDEF] has been deprecated and aliased to FILEDEF.

The most straightforward of these commands is APP HOLD, which allows you to relocate all output to a particular application. You can use this command with operations that produce output files, such as HOLD, SAVE and SAVB, as well as with CREATE SYNONYM and APP QUERY HOLD. (For details about HOLD, SAVE, and SAVB commands, see the *Creating Reports With TIBCO WebFOCUS® Language* manual.)

The APP HOLD* commands are particularly helpful when you are creating permanent files for other applications to use. However, if a command is used at an inappropriate point in the application or if it remains in effect when further steps are performed within the application, the target application may be flooded with intermediate and unintended files. Understanding the behavior of each command and the interactions among them will help you avoid this situation.

Reference: Interactions Among Output Redirection Commands

This chart describes the behavior associated with each redirection command and the interactions among them if multiple commands are used.

Command	Stand Alone	Notes
APP HOLD	Redirects all Agent output from HOLD, SAVE, SAVB, CREATE SYNONYM, and APP QUERY HOLD commands to the designated application.	When issued without a specific <i>appname</i> , APP HOLD has the effect of turning off the command.

Command	Stand Alone	Notes
<code>APP HOLDDATA</code>	Redirects the <i>data</i> from HOLD, SAVE, and SAVB operations to the designated application, but does not redirect the associated <i>metadata</i> (see Note 1 after chart).	Overrides APP HOLD.
<code>APP HOLDMETA</code>	Redirects the <i>metadata</i> from HOLD, SAVE, and SAVB operations to the designated application, but does not redirect the associated <i>data</i> (see Note 1 after chart).	Overrides APP HOLD.
<code>FILEDEF <i>ddname</i></code> <code>DYNAM ALLOC <i>ddname</i></code>	Redirects the <i>data</i> from specific HOLD, SAVE, and SAVB operations to the designated target, but does not redirect the associated <i>metadata</i> (see Note 1 after chart). The AS phrase must match the <i>ddname</i> . When there is no AS phrase, the <i>ddname</i> must match a predefined default name: HOLD for HOLD output files; SAVE for SAVE output files; and SAVB for SAVB output files.	Overrides APP HOLD and APP HOLDDATA.

Command	Stand Alone	Notes
<code>DYNAM ALLOC HOLDMAST</code>	<p>Redirects the <i>metadata</i> from HOLD, SAVE, and SAVB operations to the designated target (using the HOLDMAST <i>ddname</i>), but does not redirect the associated <i>data</i> (see Note 1 after chart).</p> <p>The recommended practice is to use this command on a request-by-request basis to avoid overriding previous output. If used as a global setting, previously held output will be overwritten with the same name.</p>	Overrides APP HOLD and APP HOLDMETA.

Note:

- ❑ Not all formats have associated metadata. For example, the HOLD FORMAT PDF command does not produce metadata, therefore, there is no metadata to redirect.
- ❑ The use of the APP HOLD command to redirect CREATE SYNONYM output is neither necessary nor desirable since the CREATE SYNONYM command directly supports application names using the syntax:

```
CREATE SYNONYM appname/synonym . . .
```

APP HOLD

The APP HOLD command defines an application in which to hold output data files (and associated Master and Access Files, if applicable) created by a HOLD, SAVE, or SAVB process in the application.

APP HOLD is intended to be used to refresh files that are common for all users of the application. It should not be used for private files since it points to an application area that is used by multiple users. If the same hold name (HOLD or AS *name*, for example) is used, conflicts between users could result.

For related information, see [Interactions Among Output Redirection Commands](#) on page 197.

Syntax: **How to Designate a Storage Location for Temporary Files**

`APP HOLD appname[/appnamea. . .]`

where:

`appname[/appnamea. . .]`

Is the application in which you wish to store output files. It can be a nested application name.

Note: Issuing APP HOLD without an *appname* turns off the effects of the command.

APP HOLDDATA

The APP HOLDDATA command designates an application as the location for storing data files created with the HOLD command. For related information, see [Interactions Among Output Redirection Commands](#) on page 197.

Syntax: **How to Designate a Storage Location for Data Files**

`APP HOLDDATA appname[/appnamea. . .]`

where:

`appname[/appnamea. . .]`

Is the name of the location for the data files created by any write process in the application. It can be a nested application name

APP HOLDMETA

The APP HOLDMETA command designates an application directory as the location for storing Master and Access Files created in the application. For related information, see [Interactions Among Output Redirection Commands](#) on page 197.

Syntax: **How to Designate a Storage Location for Master and Access Files**

`APP HOLDMETA appname[/appnamea. . .]`

where:

`appname[/appnamea. . .]`

Is the name of the location for the Master and Access Files created in the application. It can be a nested application name.

APP FI[LEDEF]

The APP FI[LEDEF] command has been deprecated and aliased to FILEDEF.

Application Metadata Commands and Catalog Metadata

Developers may want to write applications that check application metadata and decide a course of action. For example, they may want to check the existence of a file or a file date, and decide on the need for another step, such as recreation of the file. There are multiple ways to accomplish a simple check for file existence or some other attribute, that have evolved over the release history of the product. However, some of these methods have limitations. A good example of this is the STATE command, which uses a native path name for UNIX. This type of path name would not match a Windows file path and, therefore, would require IF THEN ELSE or GOTO logic to issue the correct version of the command for the operating environment, that might be quite cumbersome, depending on how often it is needed.

To solve part of this problem, commands such as STATE, FILEDEF, and DYNAM have been extended to support APP names (that is, issue APP MAP then use STATE mymap/myproc.fex). To deal with more complex issues, such as retrieving a list of available applications (APP names) and files within a particular application, a series of APP commands were developed (APP LIST and APP QUERY). However, as features such as nested applications (sub-directories) were implemented, it became apparent that a much more extended ecosystem for accessing application metadata was needed.

To satisfy this need for extended information, various internal tables were extended or created. Today the catalog/sysapps table is the primary method for accessing application metadata using standard TABLE or SELECT syntax. This is what is used in most internal applications. That is not to say that the prior methods are no longer supported. At times they can provide quick and simple coding for a specific need, but they have limitations (as noted). More complex situations require the use of the newer methods to access information. Additionally, tables such as catalog/systables and catalog/syscolum can provide additional information that is table specific, such as what DBMS a table is using and the data specification of particular columns, but they are beyond the scope of this section. It should also be noted that the newer methods occasionally overlap on how to accomplish a task. For example, a number of the catalog/sys* tables can be used to answer the question of whether a file exists. However, the tables differ from each other in the more detailed information, such as physical or application locations and attributes.

Retrieving Basic Information

The following commands return basic information about files and applications.

STATE

The STATE command allows you to check for the existence of a file. The file reference you supply can be the full path native operating system file name, or a file name prefaced with an APP name. This section only described the use of APP-name prefaced files. When an APP name is used, it does not matter if the name was natively created under the APPROOT location of the server or as an APP MAP name.

If the file does not exist, the STATE command displays a message to that effect. After issuing the STATE command, the &RETCODE system variable contains the value zero (0) if the file exists, or a non-zero value if the file does not exist.

Syntax: How to Check File Existence

```
STATE appname/filename.filetype -TYPE RETCODE &RETCODE
```

where:

appname

Is the application under which the file is stored.

filename

Is the name of the file.

filetype

Is the file type or extension of the file.

If the file exists, the &RETCODE value will be 0 (zero). Otherwise, it will be non-zero and can be used to further direct the logic of the application, typically in a -SET or a -IF command. The STATE command will also output a *not found* message. To suppress this message, use the SET TRMOUT={OFF|ON} command.

For example, the following STATE command checks the existence of the file *myproc.fex* in the *baseapp* application. The STATE command displays a message if the file does not exist. The -TYPE command displays the value zero (0) if the file exists or the value -1 if the file does not exist.

```
STATE baseapp/myproc.fex  
-TYPE RETCODE &RETCODE
```

Example: Checking the Existence of a File With the STATE Command

The following partial example suppresses the message returned by the STATE command, issues the STATE command to check if the file *myproc.fex* exists in the baseapp application, checks the return code, and creates the file if it does not exist, before continuing with the next step in the application. If the file does exist, the code immediately transfers to the next step in the application, (-RESUME label):

```
SET TRMOU=OFF
STATE baseapp/myproc.fex
SET TRMOU=ON
-IF &RETCODE EQ 0 THEN GOTO RESUME;
...
* Some code to create the file goes here
...
-RESUME
```

APP LIST

The APP LIST command alphabetically lists the applications available under the application root, APPROOT, or under an APP MAPped location. It does not care if the APP is on the current application map or not, as it is a raw list of available applications.

Syntax: How to List the Applications in APPROOT

```
APP LIST [HOLD]
```

If the HOLD option is used, the output is written to a temporary file called focappl.ftm, (FOCAPPL on PDS Deployment), which can, in turn, be used in a request to drive a report or take an action using the catalog/focappl Master File.

Limitations:

- ❑ APP LIST does not display nested application names.
- ❑ On operating systems that use case-sensitive file names (such as UNIX), uppercase physical directory names are not valid (so are not returned by APP LIST). APP names are case insensitive, but they are created on disk as lowercase, which may in turn be uppercased by the native operating system (that is, PDS uppercase file names). However, APP LIST returns them in lowercase, to be homogenous across operating systems.

Example: Using APP LIST to List and Work with Applications

The following request lists applications

```
APP LIST
```

The APP LIST output is:

```
BEGIN-APP-LIST
15/02/2000 13.36.38 baseapp
15/02/2000 13.36.38 ggdemo
15/02/2000 13.36.38 ncp
15/02/2000 13.36.38 template
END-APP-LIST
```

The following request lists applications that have been stored using the HOLD option

```
APP LIST HOLD
SQL SELECT DATE, TIME, APPNAME FROM FOCAPPL;
END
```

The APP LIST output is:

```
DATE          TIME          APPNAME
-----
15/02/2000    13.36.38      baseapp
15/02/2000    13.36.38      ggdemo
15/02/2000    13.36.38      ncp
15/02/2000    13.36.38      template
```

The following practical example of using the APP LIST HOLD command issues a TABLE request against the HOLD file to check if any files exist in the application *myapp*. If no lines are returned, the application does not exist, so it is created, and the application continues. Otherwise, the application continues without creating the application.

```
APP LIST HOLD
TABLE FILE FOCAPPL
PRINT * ON TABLE HOLD WHERE APPNAME = 'myapp'
END
-IF &LINES GT 0 THEN GOTO RESUME
APP CREATE myapp
-RESUME
```

APP QUERY

The APP QUERY command lists files within a given application. Applications and specific nested applications can be queried.

Syntax: **How to List Components**

```
APP QUERY app1[/app1a...] [app2[/app2a]...] ...
      [appn[/appna]] [HOLD]
```

where:

```
app1[/app1a...appn[/appna]]
```

Are application names. They can be nested application names. If you need to specify more application names than can fit on one line, add the continuation character (-) at the end of the first line, and continue more application names on the next line.

If the HOLD option is used, the output is written to a temporary file called focappq.ftm (FOCAPPQ on PDS Deployment), which can, in turn, be used in a request to drive a report or take an action using the catalog/focappq Master File.

Limitations: All files within an APP are listed. On systems like UNIX, this may include files of any case, so files such as MYPROC.FEX and myproc.fex may appear in a listing, but only the lowercase version would be accessed in a request.

Example: **Listing Application Files**

The following request lists application files.

```
APP QUERY abc
```

The APP QUERY output is:

```
BEGIN-APP-QUERY: abc
24/10/2014 21.38.28      4 F myproc1.fex
24/10/2014 21.38.35      4 F myproc1.fex
24/10/2014 21.37.49      4 F myappl
24/10/2014 21.32.36      0 D myapp2
END-APP-QUERY
```

The following request lists files that have been stored using the HOLD option.

```
APP QUERY ABC HOLD
SQL SELECT DATE, TIME, GMTTIME, SIZE, OTYPE, FILENAME, APPNAME FROM
FOCAPPQ ;
END
```

The APP QUERY output is:

DATE	TIME	GMTTIME	SIZE	OTYPE	FILENAME	APPNAME
----	----	-----	----	-----	-----	-----
24/10/2014	21.38.28	1414201108	4	F	myproc1.fex	abc
24/10/2014	21.38.35	1414201115	4	F	myproc2.fex	abc
24/10/2014	21.37.49	1414201069	4	D	myappl	abc
24/10/2014	21.32.36	1414200756	0	D	myapp2	abc

Note that APP QUERY ... HOLD returns a slightly extended type of information. Whitespace has selectively been removed from the above output for readability (the FILENAME column is actually 70 characters wide).

The following practical example of using the APP QUERY HOLD command checks the existence of the file *myproc1.fex* in application *abc*. If the file does not exist, the procedure exits. If the file does exist, the procedure continues.

```
APP QUERY abc HOLD
TABLE FILE FOCAPPQ
PRINT * ON TABLE HOLD
WHERE APPNAME = 'abc'
WHERE FILENAME = 'myproc1.fex'
END
-IF &LINES GT 0 THEN GOTO RESUME
-TYPE Procedure Not Found ... exiting!
-EXIT
-RESUME
```

Retrieving Extended Catalog Information

This section provides basic information about querying the server catalogs.

catalog/sysapps

The catalog/sysapps table contains metadata for physical objects on path.

This section only touches on basic uses typically needed by a developer. The Master File on disk robustly describes more attributes than are described here. You can directly study the Master File in order to understand other uses. The catalog/sys* group of files are subject to change (and are usually upwardly compatible). You should never write applications that have specific dependencies (typically on object size), which tend to cause upward compatibility issues.

***Example:* Listing Files in an APP**

The following request lists the application name, application location, file names, and file extensions in the application named *abc*.

```
TABLE FILE SYSAPPS
PRINT APPNAME APPLOC FNAME FEXT
WHERE APPNAME EQ 'abc' ;
END
```

The output (with whitespace selectively removed for readability) is:

```
APPNAME      APPLOC      FNAME      FEXT
-----      -
abc          /usr/wf/ibi/apps/abc  myproc1    fex
abc          /usr/wf/ibi/apps/abc  myproc2    fex
```

The following practical example of using the SYSAPPS table to check file existence checks the existence of the file *myproc1.fex* in the application *abc*. If it does not exist, the procedure exits. If the file does exist, the procedure transfers to the next step in order to continue:

```
TABLE FILE SYSAPPS
PRINT * ON TABLE HOLD
WHERE APPNAME = 'abc' ;
WHERE FNAME = 'myproc1' ;
WHERE FEXT = 'fex' ;
END
-IF &LINES GT 0 THEN GOTO RESUME
-TYPE Procedure Not Found ... exiting!
-EXIT
-RESUME
```

catalog/sysfiles

The catalog/sysapps table contains metadata for app name objects on a path for a select object type. The default is for file type MASTER (Master Files), but is settable for other types. Unless limited in some way, all objects (of the selected type) are displayed.

This section only touches on basic uses typically needed by a developer. The Master File on disk robustly describes more attributes than are described here. You can directly study it in order to understand other uses. The catalog/sys* group of files are subject to change (and are usually upwardly compatible). You should never write applications that have specific dependencies (typically on object size), which tend to cause upward compatibility issues.

Example: Listing APP MASTER Objects

The following request lists file names, file names with their application paths, and extensions of files with file type MASTER (the default):

```
TABLE FILE SYSFILES
PRINT FILENAME LGNAME PHNAME EXTENSION
END
```

The output (with some records and whitespace selectively removed for readability) is:

FILENAME	LGNAME	PHNAME	EXTENSION
-----	-----	-----	-----
...			
mydata	MASTER	baseapp/mydata.mas	mas
mdschema	MASTER	_edahome/catalog/mdschema.mas	mas

Example: Listing APP FOCEXEC Objects

The following request sets the file type to FOCEXEC and then prints the file names, file names with their application paths, and extensions of files with file type FOCEXEC:

```
SQL FMI SET SYSFILES FOCEXEC
TABLE FILE SYSFILES
PRINT FILENAME LGNAME PHNAME EXTENSION
END
```

The output (with some records and whitespace selectively removed for readability) is:

FILENAME	LGNAME	PHNAME	EXTENSION
-----	-----	-----	-----
...			
myproc1	FOCEXEC	baseapp/myproc1	fex
myproc2	FOCEXEC	baseapp/myproc2	fex
...			

Note: The value for LGNAME will switch to DEFAULT if the data is limited and only one object returns.

A valid value for the SQL FMI SET SYSFILES command is any valid server file type. Some examples are FOCUS, FOCEXEC, STY, PDF, or ACCESS. For a full list of valid file types, see [Designating File Types for APP Commands](#) on page 193.

Example: Using the SYSFILES Table to Check File Existence

The following practical example of using the SYSFILES table to check file existence prints the filename *myproc1* with extension *fex* (with the file type set to FOCEXEC). If no lines are returned, the file does not exist and the procedure exits. If the file exists, the procedure transfers to the point at which processing continues.

```
SQL FMI SET SYSFILES FOCEXEC
TABLE FILE SYSFILES
PRINT FILENAME ON TABLE HOLD
WHERE FILENAME = 'myproc1' ;
WHERE EXTENSION = 'fex' ;
END
-IF &LINES GT 0 THEN GOTO RESUME
-TYPE Procedure Not Found ... exiting!
-EXIT
-RESUME
```

APP HELP

The APP HELP command provides help information for all of the APP commands.

Syntax: How to Retrieve Information About APP Commands

```
APP HELP command parameters
```

where:

command

Is any valid APP command.

parameters

Are parameters that are available to or required by the command.

Restricting the Use of APP Commands

Server administrators can restrict the ability of other classes of users to change the APP environment. This setting is configurable for Application Administrators and Basic Users.

When this restriction is set, all user interfaces affected by this server setting, such as the Reporting Server browser interface, the Data Migrator desktop interface, and Managed Reporting, display only applications that are in the effective application search path, instead of displaying all applications defined by the server approot setting. In addition, the use of certain APP commands that users might otherwise issue to bypass the intended controls is restricted.

To ensure full administrative capabilities, the restriction is dynamically switched off when a user who has server administration rights logs on from client software or to the Reporting Server browser interface.

This capability is supported with all security settings.

Procedure: How to Restrict the Use of APP Commands

The server administrator can restrict the use of APP commands by other users.

1. Go to the *Access Control* page.
2. Expand the *Roles* folder for *Application Administrator* or *Basic* user.
3. Right-click a user or group and click *General Privileges*.

The *General Privileges* page opens.

4. By default, the *APATH* privilege, *Change Own Application Path (No applock)* is selected. Deselect this check box to restrict the current user or group ability to use the APP commands listed below.
5. Click *Save* to confirm the setting.

Once modification of the application path is blocked, use of the following APP commands is restricted for the designated user or group:

- The following commands generate an error message:
 - APP CREATE
 - APP DELETE
 - APP RENAME

- Although no error message is explicitly generated, the following commands are restricted if they reference an application outside of the current application path:
 - APP APPENDPATH
 - APP HOLD
 - APP HOLDDATA
 - APP HOLDMETA
 - APP MAP
 - APP PATH
 - APP PREPENDPATH
 - APP COPY
 - APP COPYF[ILE]
 - APP DELETEDF[ILE]
 - APP MOVEF[ILE]
 - APP RENAMEF[ILE]

Accessing Metadata and Procedures

Permanent files include metadata and procedures that were either created before the session by another application or remain after the session is over for use by another application.

Search Rules

Unless a file name is fully qualified with the application name, the search sequence is:

1. Current directory of the agent, which is edatemp/tsnnnnn.

2. Applications set using APP HOLDMETA for metadata files, and APP HOLDDATA for hold data files.
3. Applications set in APP PATH (including MVSAPP for z/OS).
4. The *baseapp* application.
5. The EDAHOME/catalog.
6. For stored procedures only: if the file is not found, the server checks to see if the file was allocated with a FILEDEF or DYNAM command, and if so, tries to execute it.

Example: Search Paths

The following commands follow the search path, starting with the application set by the APP HOLDMETA command:

```
APP HOLDMETA APP1
```

When a procedure is executed, and referred to by a one-part name

```
EX ABC
```

the following is executed

```
profile.fex in APP1 application
```

followed by

```
EX APP1/ABC
```

If the procedure ABC is not found in APP1, the server follows the standard search path for procedures to find and execute it.

Creation Rules for Procedure Files

Unless a file name is fully qualified or redirected to another location using an APP HOLD, APP HOLDMETA, APP HOLDDATA, FILEDEF, or DYNAM command, it is created in the temporary application area of the agent and disappears after the agent is released.

For example, on z/OS if DYNAM allocation for HOLDMAST or HOLDACC is present, the metadata files are created in the corresponding PDSs (for example, for a CREATE SYNONYM or TABLE FILE file with HOLD).

For related information, see [Output Redirection Commands](#) on page 197.

Locating Master Files and Procedures

Once your path is set, you can locate Master Files and procedures using the WHENCE command.

Syntax: **How to Locate Files**

You can issue the WHENCE command to return the fully-qualified path to the first occurrence of a file in your application path. You can issue the APP WHENCE command to return the name of the first application on your application path in which the file resides. Nested application directories are supported.

To return the location of the first occurrence of a Master File, procedure, or other FOCUS file type on your application path, issue the following command

```
WHENCE filename filetype
```

To return the name of the first application on your application path that contains a Master File, procedure, or other FOCUS file type, issue the following command

```
APP WHENCE filename filetype
```

where:

filename

Is the name of the file you are trying to locate.

filetype

Is the type of file you are trying to locate.

Example: **Locating Files**

The following command returns the location of the first occurrence of the Excel file named pivot_demo.xlsx on the application path.

```
WHENCE pivot_demo xlsx
```

The output is:

```
C:\ibi\apps\retail90\uploads\pivot_demo.xlsx
```

The following command returns the name of the application that contains the first occurrence of the Excel file named pivot_demo.xlsx on the application path.

```
APP WHENCE pivot_demo xlsx
```

The output is:

```
APPNAME: retail/uploads
```

Accessing Existing Data Files

You can allocate existing data files using the following methods:

- DATASET keyword in the Master File.
- FILEDEF command for non-FOCUS data sources (FIXED, RMS, VSAM, XML).
- USE command for FOCUS data sources.
- For the z/OS Server, native operating system services, when supported.
- DYNAM command as a USS equivalent for the MVS ALLOCATE command.
- Superseded by JCL DD card.

It is recommended that you use only one method for each allocation.

Creation Rules for Data Files

For a newly created data file, the location is determined as follows:

1. An application set by APP HOLDDATA applies to all HOLD files.
2. For FILEDEF command, one for each data file.
3. For z/OS, native operating system allocations when supported.

The request that caused the file to be created determines the file DCB parameters, such as record length, record format, and so on.

For related information, see [Output Redirection Commands](#) on page 197.

Example: Sample Allocations by JCL

The following table contains sample allocations by JCL.

VSAM	<pre>//VSAM01 DD DISP=SHR, DSN=<i>qualif</i>.DATA.VSAM</pre> <p>This type of allocation requires the <code>szero = y</code> parameter in the <code>edaserve.cfg</code> file to support sharing of BufferPool Zero.</p>
Fixed	<pre>//FIX01 DD DISP=SHR, DSN=<i>qualif</i>.FIXED.DATA</pre>
PDS	<pre>//MASTER DD DISP=SHR, DSN=<i>qualif</i>.MASTER.DATA</pre>

FOCUS	//CAR DD DISP=SHR,DSN= <i>qualif</i> .CAR.FOCUS
-------	---

Example: Sample DYNAM Commands

The following table contains samples of the DYNAM command.

VSAM	DYNAM ALLOC FILE QVASM DA <i>qualif</i> .QVSAM.VSAM SHR REUSE
Fixed	DYNAM ALLOC FILE FILE1 DA <i>qualif</i> .FILE1.DATA SHR REUSE
PDS	DYNAM ALLOC FILE MASTER DA <i>qualif</i> .MASTER.DATA SHR REUSE
FOCUS	DYNAM ALLOC FILE CAR DA <i>qualif</i> .CAR.FOCUS SHR REU

Syntax: How to Issue a FILEDEF Command

```
FI[LEDEF] filedes DISK app/[appa...]/physfile.ftm
```

where:

filedes

Is a file designation (ddname).

app/[*appa...*]

Is an application name. It can be a nested application name.

physfile.ftm

Is a physical file located in the application.

Syntax: How to Issue a FILEDEF Command to Concatenate Files

```
FI[LEDEF] concatname DISK [app1/]filename1.ext
FI[LEDEF] name2 DISK [app2/]filename2.ext
...
FI[LEDEF] namen DISK [appn/]filenamen.ext
FI[LEDEF] concatname CONCAT name2 ... namen
```

where:

concatname

Is the ddname for one of the files and the name for the concatenated files. Use this name in a request. The individual ddnames will not be available once they are used in a FILEDEF CONCAT command.

name2 ... namen

Are ddnames for the files that will be added to the concatenation.

appl ... appn

Are application names. They can be nested application names.

filename1.ext ... filenameen.ext

Are the physical file names.

Example: Concatenating Files Using FILEDEF

The following request creates three files, file1.ftm, file2.ftm, and file3.ftm.

```
APP HOLD appl
TABLE FILE WF_RETAIL_LITE
SUM COGS_US REVENUE_US
BY STATE_PROV_NAME
WHERE STATE_PROV_NAME LE 'F'
WHERE COUNTRY_NAME EQ 'United States'
ON TABLE HOLD AS file1 FORMAT ALPHA
END
-RUN
TABLE FILE WF_RETAIL_LITE
SUM COGS_US REVENUE_US
BY STATE_PROV_NAME
WHERE STATE_PROV_NAME GT 'F' AND STATE_PROV_NAME LE 'M'
WHERE COUNTRY_NAME EQ 'United States'
ON TABLE HOLD AS file2  FORMAT ALPHA
END
-RUN
TABLE FILE WF_RETAIL_LITE
SUM COGS_US REVENUE_US
BY STATE_PROV_NAME
WHERE STATE_PROV_NAME GT 'M'
WHERE COUNTRY_NAME EQ 'United States'
ON TABLE HOLD AS file3  FORMAT ALPHA
END
```

The following commands concatenate the three files.

```
FILEDEF FILE1 DISK appl/file1.ftm
FILEDEF FILE2 DISK appl/file2.ftm
FILEDEF FILE3 DISK appl/file3.ftm
FILEDEF FILE1 CONCAT FILE2 FILE3
```

The following procedure issues a request against the concatenated files.

```
TABLE FILE FILE1
SUM COGS_US REVENUE_US
BY STATE_PROV_NAME
ON TABLE SET PAGE NOPAGE
ON TABLE SET STYLE *
GRID=OFF, SIZE=8,$
END
```

The output is shown in the following image.

<u>STATE PROV NAME</u>	<u>COGS US</u>	<u>REVENUE US</u>
Alabama	\$12,774.00	\$18,612.25
Alaska	\$20,959.00	\$29,991.52
Arizona	\$9,915.00	\$15,100.19
Arkansas	\$10,822.00	\$15,000.70
California	\$61,009.00	\$83,133.90
Colorado	\$20,596.00	\$28,140.37
Connecticut	\$7,044.00	\$10,536.34
Delaware	\$5,043.00	\$6,254.51
District of Columbia	\$5,151.00	\$7,228.56
Florida	\$52,541.00	\$72,244.02
Georgia	\$27,356.00	\$37,988.80
Hawaii	\$2,770.00	\$3,997.04
Idaho	\$10,499.00	\$14,785.35
Illinois	\$43,464.00	\$60,364.26
Indiana	\$30,904.00	\$42,367.87
Iowa	\$8,964.00	\$12,316.32
Kansas	\$19,279.00	\$26,478.06
Kentucky	\$9,070.00	\$13,290.75
Louisiana	\$25,212.00	\$35,789.99
Maine	\$2,542.00	\$3,795.93
Maryland	\$23,169.00	\$32,319.96
Massachusetts	\$16,046.00	\$22,463.11
Michigan	\$45,308.00	\$62,878.30
Minnesota	\$35,815.00	\$50,591.70
Mississippi	\$6,661.00	\$10,078.13
Missouri	\$30,450.00	\$41,776.64
Montana	\$9,920.00	\$14,315.84
Nebraska	\$14,626.00	\$19,472.62
Nevada	\$10,381.00	\$14,380.39
New Hampshire	\$3,704.00	\$5,576.98
New Jersey	\$24,128.00	\$33,598.01
New Mexico	\$7,590.00	\$10,119.98
New York	\$49,956.00	\$69,800.34
North Carolina	\$27,789.00	\$39,716.65
North Dakota	\$9,478.00	\$13,552.55
Ohio	\$34,478.00	\$49,394.62
Oklahoma	\$19,010.00	\$26,164.76
Oregon	\$15,637.00	\$21,709.21
Pennsylvania	\$46,154.00	\$63,984.09
Puerto Rico	\$436.00	\$714.97
Rhode Island	\$583.00	\$818.97
South Carolina	\$4,309.00	\$7,033.86
South Dakota	\$6,684.00	\$9,311.33
Tennessee	\$19,889.00	\$27,732.33
Texas	\$101,081.00	\$141,345.92
Utah	\$7,074.00	\$10,599.14
Vermont	\$5,212.00	\$7,474.33
Virginia	\$32,148.00	\$44,582.23
Washington	\$18,101.00	\$24,042.12
West Virginia	\$5,001.00	\$7,000.03
Wisconsin	\$29,173.00	\$39,605.50
Wyoming	\$5,633.00	\$7,136.01

Syntax: How to Issue a FILEDEF Command for a Native MVS Data Set

```
FI filedes DISK "//'NATIVE.MVS.DATASET'"
```

where:

filedes

Is a file designation.

NATIVE.MVS.DATASET

Is a Native MVS data set. It can contain any number of qualifiers, up to 44 characters long.

Syntax: How to Issue a USE Command

The USE command can be issued instead of an allocation command for FOCUS data sources. The USE command is the only mechanism for accessing files on the sink machine.

Example: Sample USE Commands

The USE command supports renaming of Master Files and concatenation of data sets. The USE command is the only mechanism for accessing files on the sink machine.

Renaming a Master File

```
USE
CAR1 AS CAR
END
```

Concatenating Master Files

```
USE
CAR1 AS CAR
CAR2 AS CAR
END
```

Accessing Files on a Sink Machine

```
USE
CAR1 ON FOCUSU1
END
```

Data Set Names

If a data set name satisfies one of the following conditions, the server assumes that it is an MVS file name:

- Data set name starts with "///".

- ❑ Data set name contains no "/" and contains at least one "."

In all other cases, the name is interpreted as an HFS file name.

Syntax: **How to Define a Data Set**

The following syntax is supported:

```

DATASET=APP1/physfile.ftm
DATASET='qualif.car.data'
DATASET=qualif.car.data
    
```

In addition, on z/OS, you can use the following:

GDG files	<code>FILENAME=CARGDG,SUFFIX=FOCUS, DATASET='qualif.CARGDG.FOCUS(0)'</code>
PDS members	<code>FILENAME=CARMEMB,SUFFIX=FOCUS, DATASET=qualif.CARPDS.DATA(CARMEMB)</code>
FOCUS, VSAM, Fixed	<code>FILENAME=CAR,SUFFIX=FOCUS, DATASET=/'qualif.CAR.FOCUS'</code>

Allocating Temporary Files

Temporary files are transient files that disappear after you end a session. By default, all temporary data files (for HOLD and FOCUS files) and temporary metadata files, such as temporary Master Files and Access Files, are created in the temporary area of an agent, which corresponds to your TSCOM ID. For example, if your TSCOM ID is TS000001, your temporary files are located in //edatemp/ts000001.

For z/OS, you can control the size and location of these temporary metadata files and data files. You can specify that the temporary files reside in the hierarchical file system, MVS data sets, or in hiperspace.

Syntax: **How to Allocate Temporary Files**

To specify the allocation of your temporary files, issue the following command

```
DYNAM SET TEMP[ALLOC] {HFS|MVS|HIPER}
```

where:

HFS

Allocates temporary files to the hierarchical file system. HFS is the default value.

MVS

Allocates temporary files to MVS data sets.

HIPER

Allocates temporary files to hiperspace.

Reference: Usage Notes for Allocating Temporary Files

For z/OS, temporary metadata files can be allocated using a similar procedure to allocating permanent metadata files:

- ❑ If DYNAM allocation for HOLDMAST or HOLDACC is present, temporary files are stored in the designated PDSs.
- ❑ If DYNAM SET TEMP[ALLOC] MVS is issued; in the default temporary PDSs.
- ❑ If DYNAM SET TEMP[ALLOC] HIPER is issued; in the HIPERSPACE.

Syntax: How to Allocate Temporary Files to MVS Data Sets

To alter the default allocation parameters for temporary files for MVS data sets, issue the following command

```
DYNAM SET TEMP[ALLOC] FOR type dynam_parms
```

where:

type

Is one of the following: HOLDACC, HOLDMAST, HOLD SAVE, REBUILD, FOCUS, FOCSORT, OFFLINE, or FOC\$HOLD.

dynam_parms

Are regular DYNAM ALLOC parameters to be used as default for that type. Note that DCB parameters, if provided here, will be ignored, since they must be compatible with the file type being written.

This is similar to the functionality of IBITABLA in the SSCTL Server. The defaults should be overwritten for all cases when, in older versions, a private copy of IBITABLA existed containing different values.

Reference: System Defaults for Allocating Temporary Files to MVS Data Sets

System defaults for HOLDMAST and HOLDACC are:

```
TRKS 5 5 DSORG PO DIR 36 NEW REU
```

System defaults for all other types are:

`CYLS 5 10 DSORG PS NEW REU`

Syntax: **How to Support Long Synonym Names Using DYNAM SET LONGSYNM**

The server supports synonym names up to 64 characters. However, PDS member names cannot exceed eight characters. The server accounts for this operating environment limitation with the command DYNAM SET LONGSYNM.

A synonym comprises a Master File and, usually, an Access File. When you create a synonym with a name exceeding eight characters, the LONGSYNM setting currently in effect determines how the long name of the Master File and of the Access File will be handled.

You can issue DYNAM SET LONGSYNM anywhere SET commands are valid, including the global server profile (edasprof.prf) and a stored procedure (FOCEXEC).

DYNAM SET LONGSYNM, on servers running on z/OS, corresponds functionally to the server configuration file keyword LONGSYNM on servers running on MVS.

The syntax is

`DYNAM SET LONGSYNM {HFS | MVS | MATCH}`

where:

[HFS](#)

Specifies that each synonym whose name is longer than eight characters will be created in an HFS directory. This is the default.

[MVS](#)

Specifies that when you save a synonym with a name exceeding eight characters, the server truncates the name, preserving up to the first six characters, followed by a left curly brace ({} and a suffix number that ensures the name is unique. (The server preserves the original long name within the synonym files.)

For example, if you create a Master File named VERYLONGNAMETEST, it will be saved as VERYLO{0. If you then create a Master File named VERYLONGNAMEPROD, it will be saved as VERYLO{1.

The server chooses a suffix number by taking the next unused number in the sequence for that truncation of a Master File or Access File name. If the next number available for the Master File is different than that available for the Access File, the files will be created with different numbers. For example, if the highest Master File name truncated to VERYLO is VERYLO{8, and the highest Access File name truncated to VERYLO is VERYLO{5, and you create a synonym specifying the name VERYLONGNAMEAGAIN, the new Master File will be saved as VERYLO{9, and the new Access File will be saved as VERYLO{6.

MATCH

Works the same as the MVS setting, except that it ensures that the truncated names of a Master File and Access File synonym will always match. That is, they will be named using the same suffix number.

In the example provided for the MVS setting, if SET LONGSYNM had instead been set to MATCH, both the new Master File and the new Access File would have been named VERYLO{9.

Matching names may be a convenience for some people if they manually manage synonym files. It is less efficient than the MVS setting, however.

Temporary Space Usage and Location

The server uses temp space for the various processes. This space may be used for anything from logs and traces to transient work files (such as HOLD files). This temp space uses real disk space, but is recycled and cleaned up based on server configuration (for example, secured servers clean agent directories upon disconnect). Ultimately, a server restarts also cleans up this temporary space. If a server has been turned off, and a restart is not planned (possibly because it was a test configuration), the space may be recovered by the server admin with the following command:

```
edastart -cleardir
```

Temporary Disk Space Usage for Non-PDS Deployment

Temporary space, by default, is created under the configuration directory (ffs, wfs, or dm) in the file system used to install the product. The edatemp directory is two levels deep.

The first level directory is used for trace files if tracing is active. By default, the directory is named edatemp and is also generally referred to as the edatemp directory of EDACONF.

The subdirectory is used for temporary files and traces created by various processes, such as the TCP and HTTP listeners, the Workspace Manager, and the end-user agent processes. Each connecting user is assigned to a private data agent directory to isolate their work from other users, and that directory is cleaned up upon disconnect. The agent subdirectories are named `tsnnnnnn`, where `nnnnnn` is the ID number assigned to the data agent.

The `edatemp` directory can be configured to use a separate file system by setting the `EDATEMP` variable to point at the alternate location. The steps and syntax vary by platform, and may be done as an operating system variable or as a variable in the server configuration file (`edaserve.cfg`) before server start up.

If you are using the server configuration file, the best method is to edit the `edaserve.cfg` file from the Reporting Server browser interface *Workspace* folder, adding the line:

```
edatemp = directory
```

where

directory

Is a full path directory name, using the appropriate platform syntax. After saving, restart the server.

Note: Editing the server configuration file (`edaserve.cfg`) is the only method available for Windows.

Syntax: How to Specify the `EDATEMP` Variable

For UNIX and Linux, the `edatemp` variable can be coded in the `edastart.sh` file:

```
export EDATEMP=/u/iway/edatemp
```

Note: This also applies to IBM i and Windows.

```
DEFINE EDATEMP TMP:[TMP.EDATEMP]
```

For UNIX System Services, it can be coded in the `ISTART JCL` member under the `EDAENV dd` statement:

```
//EDAENV DD *  
EDATEMP=/u/iway/edatemp
```

For z/OS UNIX System Services only, you can set the temporary area to MVS by using the following command in the `edasprof.prfl` file or user profile.

```
DYNAM SET TEMPALLOC MVS
```

The above are global settings which affect temporary file allocations for all users.

It can be specified in the edaserve.cfg file.

```
edatemp=/u1/home/temp
```

Syntax: How to Pre-Allocate Temporary Files

You can pre-allocate an individual file for a user, using the following techniques:

❑ For UNIX and Linux:

```
FILEDEF XXX DISK /u/another/area/xxx.dat
```

where:

```
/u/another/area
```

Has enough free space to hold the file.

❑ For Windows:

```
FILEDEF XXX DISK C:\tmp\another\area\xxx.dat
```

where:

```
tmp\another\area
```

Has enough free space to hold the file.

❑ For z/OS UNIX System Services:

You can use the same FILEDEF syntax indicated for UNIX, or you can use the DYNAM command to direct the temporary file to MVS.

```
DYNAM ALLOC FILE XXX SPACE
```

Note that, for the DYNAM command, you must specify the amount of space required.

❑ For FOCUS files:

You can use the FILEDEF and USE commands to create a FOCUS file outside the edatemp area.

```
FILEDEF NAME DISK /{pathname}{filename}.foc .....
USE NAME NEW
END
```

Syntax: How to Dynamically Allocate FOCUS Files on z/OS

You can dynamically allocate FOCUS files on z/OS with the USE command. The command is

```
DYNAM ALLOC FILE ddname SPACE  
USE ddname AS masterfile  
END
```

where:

ddname

Is the DDNAME.

masterfile

Is the Master File name.

If the DDNAME and Master File name are the same, use just the command:

```
DYNAM ALLOC
```

Application Tools

An application program may be as simple as a TABLE FILE ... END request or as elaborate as a series of replaceable variables (known as amper variables) controlled by dynamic inputs, GOTO statements, and/or looping logic (known collectively as Dialogue Manager). While Dialogue Manager is fully detailed in the *TIBCO WebFOCUS® Developing Reporting Applications* manual, a limited set of these commands and tools are also noted here due to their integral use in APP management.

EX Procedure and Amper Variables

The EX procedure executes another procedure, as found on the application path. The procedure extension (on non-PDS platforms) is not required, and is assumed to be .fex. The full syntax is:

```
EX[EC] [appname/]procedure[.fex] [parameter=value[,parameter=value][,...]
```

Parameters may be names (such as LASTNAME) or numeric (such as 1), and may be mixed within the list. A specific APP name can be referenced and does not require the APP to be on the application path prior to execution. Numerics can also be assumed positionally, as in this example:

```
EX mytest DETAIL=ALL,QUARTERLY,PERSON=JAMES MADISON,5=MANAGER
```

and the procedure would contain amper variables (parameters) for &DETAIL, &2, and &PERSON, and would use the context of:

```
TABLE FILE ...  
WHERE PERSON IS '&PERSON' AND PERIOD EQ '&2'  
...  
END
```

There are two types of amper variables, single (local) and double (global). Single amper variables are only active within the procedure they are created in. Double amper variables are good for the life of the session. Additionally, there are predefined, single amper variables, such as &LINES (records returned from a request), &FOCUSUSER (current user ID), &FOCFOCEXEC (name of the current executing procedure), &FOCCODEPAGE (the server NLS code page), and many others.

EX EDAMAIL

EDAMAIL is an internal procedure that allows you to send emails from a procedure with the content of a specific file included in the email body or as an attachment. This is useful for sending email alerts for events ranging from special error conditions to simple report completion. EDAMAIL requires a configured (and working) email SMTP Server node for the server workspace.

There are two forms of syntax for EDAMAIL:

- The original form uses positional parameters, and is limited to the original specification.
- The extended form uses non-positional, name-value pair parameters. The extended form supports additional SMTP email client features such as Reply To and Importance.

The extended form is assumed if the first parameter contains an equal sign (=) indicating a name-value pair is being passed.

The basic form of email address (that is, `foo@foo.com` or `<foo@foo.com>`) may be used in either the original or the extended form of EDAMAIL. Extended forms of email addressing (such as `My Support <support@foo.com>`, `"My Support" <support@foo.com>`, `"My Support <support@foo.com>"`, and `"Support, Me <support@foo.com>"`) are only supported in the extended form of EDAMAIL.

Multiple addresses may be use for parameters that support multiple addresses (for example, `to`) when properly formed and delimited.

For multiple addresses and basic EDAMAIL, the separator must be a semi-colon (;), and the overall string may be, optionally, enclosed in double quotation marks, but remember that basic EDAMAIL only supports simple addresses.

For multiple addresses and extended EDAMAIL, the separator is normally a semi-colon (with exceptions where a comma may be used), plus the overall string must be enclosed in single quotation marks if any of the addresses contains a comma (for example, `Support, Me <support@foo.com>`). A comma (,) may be used as a separator if the email addresses are all simple addresses (that is, `foo@foo.com` or `<foo@foo.com>`), plus the overall string must be enclosed in single quotation marks (') or double quotation marks (").

The actual use and display of the email headers created by EDAMAIL are a function of the SMTP Server in use (and any intervening SMTP hops), plus the user email client (such as Outlook, Gmail, Hotmail, and so on). They are not directly controlled by the sender email headers. Therefore, an email may not always exhibit the expected behavior in all email client environments. For example, older email clients may not recognize newer header types, and may ignore Reply To. Any functional issues should first be investigated by an experienced SMTP/Email administrator to determine if they are client-related.

The EDAPRINT log records all EDAMAIL executions. It shows the user ID that executed the EDAMAIL command and sender and addressee information. For example:

```
02/01/2019 16:58:41.395 I EDAMAIL u=PTH
\srvadmin,from=user2@ibi.com(),to=(user1@ibi.com)
```

Syntax: How to Use EX EDAMAIL in the Positional Parameter Form

The syntax is:

```
EX EDAMAIL to, cc, bcc, from, subject, [flag], filetype, data
```

where:

to

Is the email recipient. Multiple addresses are allowed, using a semi-colon (;) as the address separator.

cc

Is the carbon copy recipient. Multiple addresses are allowed, using a semi-colon (;) as the address separator.

bcc

Is the blind carbon copy recipient. Multiple addresses are allowed, using a semi-colon (;) as the address separator.

from

Is the email sender.

subject

Is the email subject string. If the subject string contains a comma (,), the string must be enclosed in single quotation marks (') or double quotation marks (").

flag

If set to a or A, the file is sent as an attachment. Otherwise, it is included as the body of the email. All other values are ignored. The value is also ignored if the filetype parameter is blank.

filetype

Defines the data file type for an email message body that uses a file. Any application file type is valid, including MASTER, FOCEXEC, HTML, TEXT, and so on. Leave the parameter empty to use the inline email message body feature.

data

Is the inline data stream for the email message body or the [app/]filename file containing the email message body. The inline data stream feature requires an empty filetype parameter. The EDAMAIL feature can also spread an inline email message body on to multiple lines using the -LINES {n|*} feature.

If an inline data stream message body is spread across multiple lines in the procedure, the resulting email is a single line of output. Multi-line message bodies are respected when the message body from a file option is used.

Example: Mailing an HTML File as Message Body

```
TABLE FILE file1
PRINT A B C
ON TABLE HOLD AS MYFILE FORMAT HTML
END
EX EDAMAIL user1@corp1.com, user2@corp1.com, File1 Report,,HTML, MYFILE
```

Example: Mailing an HTML File as a Message Attachment

```
TABLE FILE file1
PRINT A B C
ON TABLE HOLD AS MYFILE FORMAT HTML
END
EX EDAMAIL user1@corp1.com, user2@corp1.com, File1 Report,a,HTML, MYFILE
```

Example: Mailing a Multi-line Inline Message

```
...
EX -LINES * EDAMAIL user1@corp1.com, user2@corp1.com, &SUBJECT,,,
    Run result for &TESTNAME is:
&RESULT
EDAMAIL*
```

Syntax: **How to Use EX EDAMAIL in the Extended Form**

The EDAMAIL extended form using name-value pairs is activated by the detection of an equal sign (=) in the first parameter. Parameter names are not case-sensitive and may be in any order, but the message parameter (if used) must be last.

The syntax is:

```
EX EDAMAIL to=addresslist,
[toaddr= {addresslist|[%[app/]addresslist.fex}, ]
[cc=cadrlist1,]
[ccaddr=cadrlist2,]
[bccaddr=bcadrlist,] [from=address,]
[fromaddr=address,] [replyto=address,] [importance=low/normal/high,]
[subject=string,] [flags=value,] [filetype=extension,]
[filename=file,] [message=body message]
```

where:

to=addresslist

Is the email recipient displayed by the user email client (and used in an email client Reply To All, if a toaddr header is not supplied). Multiple addresses are allowed, using a semi-colon (;) as the address separator. A comma (,) is also allowed as a separator, if the overall address string is enclosed in single quotation marks (').

If the to parameter is used in conjunction with toaddr, the value of to may be an arbitrary string, such as *Group Managers*, which most email clients will display as a pseudo title in the To field, and will not display the actual addresses used in the toaddr parameter. A forced blank can be supplied for the To field by using *to=""*.

toaddr=addresslist

Is the email recipient. If not supplied, SMTP servers will use the to header. Email clients will use the toaddr value in an email client Reply To All. Multiple addresses are allowed, using a semi-colon (;) as the address separator. A comma (,) is also allowed as a separator, if the overall address string is enclosed in single quotation marks (').

toaddr=%[app/]addresslist.fex

Is a FOCEXEC that generates an email recipient list at run time. Email clients will use the toaddr value in an email client Reply To All. The actual FOCEXEC may go against any data source for the addresses, but must PRINT a single column and use ON TABLE PCHOLD FORMAT COMT FORMATTED syntax to format the data (as one email address enclosed in double quotation marks (") per row).

The actual FOCEXEC extension must be .fex in the specification. This feature also only works in a running server context. It is not supported in any other mode.

cc=cadrlist1

Is the carbon copy recipient displayed by the user email client (and used in an email client Reply To All, if a ccaddr header is not supplied). Multiple addresses are allowed, using a semi-colon (;) as the address separator. A comma (,) is also allowed as a separator, if the overall address string is enclosed in single quotation marks (').

If the cc parameter is used in conjunction with ccaddr, the value of cc may be an arbitrary string, such as *Group Managers*, which most email clients will display as a pseudo title in the Cc field, without displaying the actual addresses used in the ccaddr parameter. A forced blank can be supplied for the Cc field by using `to=""`.

ccaddr=cadrlist2

Is the carbon copy recipient. If not supplied, SMTP servers will use the cc header. Email clients will use the ccaddr value in an email client Reply To All. Multiple addresses are allowed, using a semi-colon (;) as the address separator. A comma (,) is also allowed as a separator, if the overall address string is enclosed in single quotation marks (').

bccaddr=bcadrlist

Is the blind carbon copy recipient. Email clients will use the bccaddr value in an email client Reply To All. Multiple addresses are allowed, using a semi-colon (;) as the address separator. A comma (,) is also allowed as a separator, if the overall address string is enclosed in single quotation marks (').

from=address

Is the email sender displayed by the email client (and used in an email client Reply, unless overridden by a fromaddr or Reply To header). If the from parameter is used in conjunction with fromaddr, the from value may be an arbitrary string, such as *The Boss*, which most email clients will display as a pseudo title in the From field, and will not display the actual address used in the fromaddr parameter.

fromaddr=address

Is the email sender. If not supplied, most email clients will use the From header when doing a reply.

replyto=address

Is the email sender. If not supplied, most email clients will use the fromaddr or from parameter value.

importance=low/normal/high

Is the email importance level for email clients that support importance flags. Valid values are high, normal, and low.

subject=string

Is the email subject string. If the subject string contains a comma, the string must be enclosed in single or double quotation marks.

flags=value

If set to a or A, the file is sent as an attachment. Otherwise, it is included as the body of the email.

filetype=extension

Defines the data file type for an email message body that uses a file as the body. Any application file type is valid, including MASTER, FOCEXEC, HTML, TEXT, and so on. Leave the parameter out to use the inline email message body feature.

filename=file

Defines the data file for an email message body that uses a file as the body. Leave the parameter out to use the inline email message body feature.

message=body message

Is the inline data stream containing the email message body. If used, it must be the last parameter in the EDAMAIL command. To use the inline data stream feature, the filetype and filename parameters cannot be supplied. The data stream may also be spread onto multiple lines if EDAMAIL is used with the EX -LINES {n|*} feature.

If an inline data stream message body is spread across multiple lines in the procedure, the resulting email is a single line of output. Multi-line message bodies are respected when the message body from a file option is used.

Example: Mailing an HTML File as a Message Body Using Multiple Addresses and Extended Form

```
TABLE FILE file1
PRINT A B C
ON TABLE HOLD AS MYFILE FORMAT HTML
END
EX EDAMAIL to=Managers,toaddr=user1@corp1.com;user2@corp1.com,
from=support1@corp1.com,subject=File1 Report, filetype=HTML,
filename=MYFILE
```

Managing Adapters


Adapters must be configured to allow a WebFOCUS® Reporting Server to access data in different types of databases. You must add and configure an adapter for every database type you want to access.


This topic describes briefly how to declare connection attributes and authentication information, when required, using the Reporting Server browser interface or the Data Migrator desktop interface.

In this chapter:

- [Changing Common Adapter Settings](#)
 - [Configuring an Adapter](#)
 - [Changing the Adapter Configuration](#)
 - [Configuring a Remote TIBCO WebFOCUS Reporting Server](#)
 - [Complete Adapter List](#)
 - [SQL Automatic Passthru Status](#)
 - [Determining SQL Optimization](#)
 - [Data Type Support Report](#)
 - [Displaying License Counts](#)
-

Changing Common Adapter Settings

On the Get Data page, click *Settings* ().

On the Advanced Mode page, click *Options* (), then click *Common Settings* to customize data type mappings and other miscellaneous settings that are common to multiple adapters.

Adapter-specific settings are customizable as well, and can be accessed by right-clicking a configured adapter and selecting *Change Settings* from the context menu.

Configuring an Adapter

You can configure any supported adapter for use with a WebFOCUS Reporting Server from the Reporting Server browser interface or the Data Migrator desktop interface.

If an adapter has already been configured, it is listed in the Server Files list of the Get Data page. Adapters that have not been configured can be configured by clicking the plus sign (+) on the action bar. Adapters are grouped as follows:

- SQL
- XML Based
- ERP
- OLAP
- Procedures
- Statistics
- GIS
- Files
- DBMS
- Social Media
- SCCS
- Search Engines
- Remote Servers

For general instructions about adding an adapter and declaring connection attributes, see [How to Declare Connection Attributes](#) on page 235.

Note: Although all database configuration panes have the same look and feel, the parameters are specific to each database. To obtain *detailed* configuration information for an adapter, click the ? next to each parameter.

Most of the adapter configuration information is contained in two configuration files, EDASPROF.PRF and EDASERVE.CFG.

- EDASPROF.PRF, or the global profile, contains data adapter connection information.
- EDASERVE.CFG contains database release and access method information.

Your entries on the Reporting Server browser interface or the Data Migrator desktop interface populate these files.

Procedure: How to Declare Connection Attributes

You can configure the adapter from either the Reporting Server browser interface or the Data Migrator desktop interface.

1. From the Reporting Server browser interface Applications page, click *Get Data*.

or

From the Data Migrator desktop interface, expand the *Adapters* folder. The Adapters folder opens. Expand the *Available* folder, if it is not already expanded.

From the Reporting Server browser interface, click the New Datasource button (+). The list of available adapters opens. There is a drop-down menu you can use to show adapters in a specific category. You can also search for an adapter by clicking the search icon (magnifying glass).

Note: Prerequisites may be required for the adapter you want to configure, such as client software that needs to be installed or system variables that must be set. For information about prerequisites for configuring the adapter, right-click the adapter name on the New Datasource list, and click *Prerequisites* on the context menu. After an adapter is configured, the *Prerequisites* link is available on the context menu when you right-click an adapter name on the Configured Adapters list.

2. Right-click the adapter name and/or version and click *Configure*.

The Add Adapter to Configuration pane opens.

3. Enter values for the parameters required by the adapter, as described in the *TIBCO WebFOCUS® Adapter Administration* manual.
4. Click *Configure*. The configured adapter is added to the Adapters list in the resources tree.

Once an adapter is configured, you can right-click the adapter name or connection to perform a variety of other tasks, including creating synonyms, testing or deleting the connection, and changing or viewing the connection properties. Specific options vary by adapter. For information, see the *TIBCO WebFOCUS® Adapter Administration* manual.

Procedure: How to Test an Adapter Connection

1. From the Reporting Server browser interface Application page, click *Get Data*.

or

from the Data Migrator desktop interface, expand the *Adapters* folder. The Adapters folder opens.

Expand an adapter.

2. Right-click the adapter connection and, if available for the selected adapter, click *Test*.

The test results appear in the right pane. Test results vary, depending on adapter.

Changing the Adapter Configuration

The *Configured Adapters* folder or pane lists the adapters that have already been configured for the current operational instance.

The following general information applies.

- Right-click an Adapter folder in the navigation pane and choose from the tasks that are available for that adapter. The following table lists and describes the available tasks, or functions.

Function	Description
Add connection	Adds a connection for adapters that have the functionality to support multiple connections. These can be created in any level profile.
Change Settings	Displays the Change Setting page, where you can change a variety of settings in effect for the selected adapter.
Remove	Removes the adapter configurations and all connection settings from the list of configured adapters.
Help	Provides specific information for the adapter.

- Open the Adapter folder in the navigation pane and right-click a version or connection to choose from the tasks that are available for the selected adapter. The following table lists and describes the available tasks or functions.

Function	Description
Show DBMS objects (for different adapters, the working for this option can vary)	Opens the first Create Synonym pane, where you can begin to enter the information required to create a synonym. These panes vary for different adapters. For information, click the ? next to each parameter.

Function	Description
Properties	Enables you to view and change connection parameters. Note: Duplicate entries are not verified.
Test	Runs a select statement for name/owner against a catalog table with a limit of up to 15 records. See How to Test an Adapter Connection on page 235.
Duplicate Connection	Creates a new connection with the same parameters and a different connection name.
Test Bulk	Issues a report on whether bulk load is supported for this adapter.
Test DBMS Case Sensitivity	Issues a report on case sensitivity for this adapter.
Impact Analysis	Produces a report of the synonyms that use the connection.
Delete/Remove	Deletes connection settings. Note: Entries are deleted from the profile in which they were created.

Syntax: How to Change the Connection Attributes for an Adapter From the Command Line

The UPDATE_PROFILE option of the edastart -x command enables you to change adapter connection attributes, including the password for an adapter connection. The user must have the ADPTP privilege to execute it. The syntax is

```
UPDATE_PROFILE profile SUFFIX suffix CONNECTION conname NEW_PASS password
```

where:

profile

Is the name of the profile in which to store the connection attributes, for example, edasprof.

Any profile (edasprof.prf , user/group/role profile) can be updated using this command.

suffix

Is the suffix value for the adapter.

conname

Is the adapter connection to be updated.

password

Is the new password.

Note: The password will be encrypted when it is written to the profile.

The following command updates the password for connection CON01 for the Adapter for Oracle to *NewPasswd1* and stores the connection attributes in the *edasprof.prf* file.

```
edastart -x "EX UPDATE_PROFILE edasprof SUFFIX SQLORA CONNDECTION CON01  
NEW_PASS NewPasswd1"
```

The following command updates the password for connection CON02 for the Adapter for SQL Server to *NewPasswd2* and stores the connection attributes in *user1.prf*.

```
edastart -x "EX UPDATE_PROFILE user1 SUFFIX SQLMSS CONNECTION CON02  
NEW_PASS NewPasswd2"
```

The keywords in the command can be replaced by commas. For example:

```
edastart -x "EX UPDATE edasprof,SQLORA,CON01,NewPasswd1"
```

Configuring a Remote TIBCO WebFOCUS Reporting Server

Configuring remote servers allow you to add nodes to the communications configuration. The connection attributes for a remote server are written to the server, user, group, or role profile, which can be chosen during configuration.

Remote servers can be configured from the navigation pane of the Adapters page. Right-click *Remote Servers* and select *Configure*. The Remote Server Configuration page opens. For configuration information, click the ? next to each parameter.

This adapter allows applications to access data sources that reside on a remote server.

Complete Adapter List

You can generate a list of every adapter with accompanying information about its major features using the Reporting Server browser interface.

How to Display the Complete Adapter List Report

1. Go to the *Get Data* page.
2. From the Information menu of the page or the Options menu of the Advanced Mode page, click *Complete Adapter List*.

The Complete Adapter List Report is displayed, as shown in the following image.



Category	Sub-category	Adapter	Bulk Load	SQL Merge	SQL SP	CDC API	Unicode	Vendor Notes
SQL	1	Db2/Db2 Warehouse	Standard	Full Functionality	Yes	ZOS IFU/LUW ReadLog/IS Log API	Opt	Also for Big SQL
SQL	1	MS Azure Synapse Analytics (formerly SQL DW)	Extended: BCP	Append Only	Yes		Opt	
SQL	1	MS SQL Server ODBC/AzureDB	Extended: BCP	Full Functionality	Yes	CDC functions API	Opt	
SQL	1	MS SQL Server OLE DB/AzureDB	Extended: BCP	Full Functionality	Yes	CDC functions API	Opt	Also for PolyBase
SQL	1	MySQL	Extended: LOAD DATA	Append Only	Yes	MYSQBINLOG Utility	Opt	
SQL	1	Oracle	Extended: LOAD API	Full Functionality	Yes	Log Miner API	Opt	
SQL	1	Teradata	Extended: TPT API	Full Functionality	Yes		Opt	
SQL	2	1010data	Extended: POWER LOADER		Yes			
SQL	2	Amazon Athena	Extended: Amazon AWS S3 MV	Append Only			Opt	
SQL	2	Amazon Redshift	Extended: COPY	Append Only			Opt	
SQL	2	Apache Drill	Extended: HADOOP FS				Opt	MapR & Apache Drill
SQL	2	Apache Hive	Extended: LOAD DATA	Full Functionality			Opt	Hortonworks, MapR & Apache Hiveserver2
SQL	2	Apache Phoenix		Append Only			Opt	Apache HBase & MapR-DB
SQL	2	Apache Spark SQL	Extended: LOAD DATA				Opt	Spark SQL
SQL	2	Cloudera Impala	Extended: HADOOP FS	Append Only			Opt	Cloudera Impala
SQL	2	Couchbase (Pilot)	Extended: cbimport	Append Only			Opt	
SQL	2	ESP		Append Only			Opt	ESP

The report lists the category, subcategory, and name of each adapter. It also lists its bulk load type, whether it supports SQL Merge, SQL SP, the change data capture API, and Unicode support. There is also a column for vendor notes.

SQL Automatic Passthru Status

You can find the Automatic Passthru (APT) status for a function or keyword from a report available from the Reporting Server browser interface.

How to Display the SQL APT Status Report

1. Go to the *Get Data* page.
2. From the *Information* menu, or the *Options* menu of the Advanced Mode page, click *SQL APT Status*.

The SQL APT Status Report is displayed, as shown in the following image.

SQL APT Status Report

Type	Name	Adapter Subcategory										
		Database Name	DB2	MS SQL Server (OLE DB obsolescent)	MS SQL Server ODBC	MySQL	Oracle	Teradata	1010data	Apache Drill	Apache Hive	A
function	ABS(arg)		Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Convert	Pass as is	C
	CAST(arg AS data_type[(length[, scale]])		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	CEIL(number)		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	CHAR(arg[, date_format])		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	CHAR_LENGTH(source_string)		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	COALESCE(arg1, arg2[, ... argn])		Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	Pass as is	P
	CONCAT(arg1, arg2)		Pass as is	Convert	Convert	Convert	Pass as is	Convert	Convert	Convert	Convert	C
	COUNTBY(arg)		Convert	Convert	Convert	FOCUS	Convert	Convert	FOCUS	Convert	Convert	C
	CURRENT_DATE()		Convert	Convert	Convert	Convert	Convert	Convert	FOCUS	Convert	Convert	C
	CURRENT_TIME([precision])		Convert	Convert	Convert	Pass as is	FOCUS	Pass as is	FOCUS	Convert	Convert	C
	CURRENT_TIMESTAMP([precision])		Convert	Convert	Convert	Convert	Convert	Convert	Pass as is	FOCUS	Convert	C
	DATE(arg)		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	DAY(arg)		Pass as is	Pass as is	Pass as is	Pass as is	Convert	Convert	Convert	Convert	Pass as is	C
	DAYS(arg)		Pass as is	Convert	Convert	Convert	Convert	Convert	FOCUS	FOCUS	Convert	F
	DAY_OF_YEAR		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C
	DECIMAL(arg[, length[, dec_places]])		Pass as is	Convert	Convert	Convert	Convert	Convert	FOCUS	Convert	Convert	C
	DIGITS(arg[, num_digits])		Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	Convert	C

The report lists the APT status for each function and keyword by each SQL adapter. The possibilities are:

- Pass as is.** Do not convert, pass to the DBMS as is.
- Convert.** Convert into a DBMS-specific implementation and pass to the DBMS.
- FOCUS.** Force the SQL to be processed by FOCUS.

Determining SQL Optimization

You can determine which functions are optimized to their SQL counterparts by adapter in the SQL Optimization Report available from the Reporting Server browser interface.

How to Display the SQL Optimization Report

1. Go to the *Get Data* page.
2. From the *Information* menu, or the *Options* menu of the *Advanced Mode* page, click *SQL Optimization Report*.

The Filter SQL Optimization Report page opens, as shown in the following image.

Select Adapter Subcategory

Select Function Category

? **Filter by function name AND/OR description**

Show Function Description

Show DBCS Configuration

Show key parameter value(s)

Show Report in PDF

3. Optionally, select an *Adapter Subcategory* from the corresponding list box.
4. Optionally, select a *Function Category* from the corresponding list box.
5. Optionally, click *Filter by function name AND/OR description* to open text boxes for entering a string or pattern for matching function names and/or a string for matching function descriptions.
6. Optionally, select the *Show Function Description*, *Show DBCS Configuration*, or *Show key parameter value(s)* check boxes.
7. Click *Show Report*.

The SQL Optimization Report is displayed, as shown in the following image.

SQL Optimization Report

Adapter Subcategory	Database Name	Db2/Db2 Warehouse	MS Azure Synapse Analytic	MS SQL Server ODBC/AzureD	MS SQL Server OLE DB/Azur	MySQL	Oracle	Teradata	1010data	Amazon Athena	Amazon Redshift
Function	Function Type										
Function category: Aggregation Operation											
PCT.CNT.	Numeric	--	--	--	--	--	--	--	--	--	--
PCT.	Numeric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AVE.DST.	Numeric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ASQ.	Numeric	--	--	--	--	--	--	--	--	--	--
AVE.	Numeric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MDN.	Numeric	--	--	--	--	--	--	--	--	--	--
MDE.	Numeric	--	--	--	--	--	--	--	--	--	--
CNT.DST.	Integer	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CNT.	Integer	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FST.	derived	--	--	--	--	--	--	--	--	--	--
LST.	derived	--	--	--	--	--	--	--	--	--	--
MAX.	derived	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MIN.	derived	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
STDP.	Numeric	--	--	--	--	--	--	--	--	--	--
STDS.	Numeric	--	--	--	--	--	--	--	--	--	--
SUM.DST.	Numeric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Data Type Support Report

SQL Data Type mapping options are available in a report available from the Reporting Server browser interface.

How to Access the Data Type Report

1. Go to the *Get Data* page.
2. From the *Information* menu, or the *Options* menu of the Advanced Mode page, click *Data Types*.

The Filter Data Types Report page opens, as shown in the following image.

The screenshot shows a web interface for filtering data type reports. It includes three dropdown menus: 'Adapter Subcategory' (set to 'All'), 'Adapter', and 'Server Data Types'. Below these is a checkbox labeled 'Show Report in PDF' which is currently unchecked. At the bottom of the filter section is a 'Show Report' button.

3. Select an *Adapter Subcategory* from the corresponding list box.
4. Select an *Adapter* from the corresponding list box.
5. Select a *Server Data Type* from the corresponding list box.

6. Click *Show Report*.

The Data Types Report is displayed, as shown in the following image.

Data Type Category	Vendor Data Types	Type Range	Server USAGE	Server ACTUAL	Server DDL	Remarks
DB2, all platforms						
Date-Time	DATE		YYMD	DATE		
	TIME		HHIS	HHIS		
	TIMESTAMP		HYYMDm	HYYMDm		
LOB and Other	CLOB		TX50	TX		
	BLOB		BLOB	BLOB		
	DATALINK		N/A	N/A		
	XML		TX50L	TX		
Numeric	SMALLINT		I6	I2		
	INTEGER		I11	I4		
	BIGINT		P20	P10		
	DECIMAL(p,s)/NUMERIC(p,s)	p=1..31,s=0	Pn	Pk		n=p+1,k=(p/2)+1
	DECIMAL(p,s)/NUMERIC(p,s)	p=1..31,s>0	Pn.m	Pk		n=p+2,m=min(s,31),k=(p/2)+1
	REAL		F9.2	F4		
	DOUBLE		D20.2	D8		

Note: You can also display the report as a PDF by selecting the *Show Report in PDF* check box.

Displaying License Counts

You can display license counts from the Reporting Server browser interface.

Go to the *Get Data* page.

From the *Information* menu, or the *Options* menu on the Advanced Mode page, click *Adapters License Count*.

A report displays showing your license count, the number of adapters you are licensed for, the number of licensed adapters that are configured, and the number of free adapters that are configured.

Metadata

The data source being used determines the type of metadata that is required. For example,

- ❑ When the server accesses a relational data source, it needs to know how to interpret the data stored there. You must create a synonym that describes the structure of the data source and the server mapping of the data types.
- ❑ When the server invokes a transaction or procedure, it needs to know how to build the request, what parameters to pass, and how to format an answer set from the response. You must create a synonym that describes the layout of the request/response area.

Whatever your data source, the adapter you are using manages the synonym creation process for you, creating a synonym that meets your specific requirements.

This topic describes briefly how to create, test, and otherwise manage a synonym, using the Reporting Server browser interface or the Data Migrator desktop interface.

In this chapter:

- ❑ [Creating Synonyms](#)
 - ❑ [Testing Synonyms](#)
 - ❑ [Managing Synonyms](#)
 - ❑ [LNGPREP Utility: Preparing Metadata Language Files](#)
 - ❑ [Uploading Files](#)
 - ❑ [Getting Data](#)
 - ❑ [Using Data Assist](#)
-

Creating Synonyms

Synonyms define unique names (or aliases) for each table, view, transaction or procedure that is accessible from the server. Synonyms are useful because they hide the underlying data source location, identity, transaction or procedure from client applications. They also provide support for extended metadata features of the server, such as virtual fields and additional security mechanisms.

You can create synonyms for supported adapters from the Reporting Server browser interface or the Data Migrator desktop interface.

Note: Although all synonym creation panes have the same look and feel, the parameters are specific to each adapter. To obtain detailed information for an adapter, click a ? next to a parameter.

When you use Get Data to upload files, a Data Flow to load data to a target, or Quick Copy to copy data from one DBMS to another, the column names generated may be truncated in order to account for variations in DBMS-specific name length limits.

The name length limit used has been increased to the maximum length supported by the DBMS. In most cases this is 128 characters, although some have shorter limits, including the following:

- SAP HANA, Amazon Redshift.** 127 bytes.
- MySQL, MariaDB.** 64 bytes.
- PostgreSQL, Greenplum, Hyperstage.** 63 bytes.
- Progress.** 32 bytes.
- Oracle, DB2 on z/OS.** 30 bytes.

Procedure: How to Create a Synonym

To create a synonym, you must have previously configured the adapter. You can create a synonym from the Applications or Adapters pages of the Reporting Server browser interface.

1. Go to the *Applications* page.
2. Click *Get Data*.

The *Get Data* page opens. You can create a synonym from the *Get Data Simple Mode* page or the *Advanced Mode* page. *Advanced Mode* gives you the option of editing the *Business View* prior to saving the synonym.

By default, *Get Data* opens in *Simple Mode*. If you want to open in *Advanced Mode* automatically, do the following:

1. Navigate to the *Workspace* page.
2. Click *Settings*, then click *Settings for Reporting Server browser interface Preferences*.

The *Change Settings for Reporting Server browser interface Preferences* page opens.

- Expand the *Data Assist* section and scroll down, as shown in the following image.

The screenshot shows a configuration page for Data Assist with the following parameters and values:

- NTM_SUPPORT_BV_PLUS_FOR_CUBE**: Yes
- SUPPORT_BIZ_FOLDERS_FOR_CUBE**: Support Business Folders for Cube
- AUTORUN_COMPOUND**: Run on Change
- AUTO_RUN_FOR_COMPOUND_LAYOUT**: Auto-run for Compound/Layout
- USE_MONOTREE**: Yes
- USE_MONOTREE_DESCRIPTION**: Use monotree when selecting files in the Data Assist (Join, Comp)
- AUTO_REDO_PREPARE**: No
- AUTOMATICALLY_REDO_PREPARE**: Automatically redo prepare
- NTM_DEFAULT_BV_INCLUDE_KEYS**: No
- INCLUDE_KEYS_FOR_DEFAULT_BV**: Include Keys for Default BV
- GET_DATA_ADV_DEFAULT**: No

At the bottom of the form, there are three buttons: "Reset to defaults", "Cancel", and "Save".

- Select Yes for *GET_DATA_ADV_DEFAULT*.
- Click Save.
- Right-click a connection for a configured adapter and click *Show DBMS objects*.
If the connection is not for a DBMS adapter, the corresponding menu option may be *Show files* or some other item that indicates that you will select objects for which to create synonyms.

The synonym creation page opens.

- Enter values for the parameters required by the adapter as described in the synonym creation parameters reference.
- After entering the parameter values, click *Add*. It may be *Next*, or *Create Base Synonyms*, or some other option that identifies what is being created.

The Status page lists the synonyms in the application directory you selected.

If you are creating the synonym from the advanced mode page, you can edit the Business View before saving the synonym.

Note:

- ❑ Right-clicking an application in the Data Migrator desktop interface also provides a *New/Synonym via Synonym Editor* option, which enables you to create a synonym using that tool. For more information, see the *TIBCO® Data Migrator User's Guide*.

Reference: Checking Synonyms

Once you have created a synonym, you can view its properties by right-clicking the synonym and selecting *Properties* from the context menu.

The Properties page opens and shows the location, adapter, date last modified, and privileges of the connected user for this synonym.

Clicking the *Check File* button displays the number of errors, number of segments, number of fields, and length of all fields for the synonym.

Testing Synonyms

After you have created a synonym, it is good practice to ensure that it can retrieve data from a data source. You can quickly test this functionality on a sample of up to 50 rows.

You can test a synonym from either the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page.

or

from the Data Migrator desktop interface, expand the Server node folder.

2. Expand an application folder.
3. Right-click a synonym and select *Sample Data*.

The sample data is displayed.

Managing Synonyms

Once you have created a synonym, you can right-click the synonym name in the navigation pane of either the Reporting Server browser interface or the Data Migrator desktop interface to access the following options:

Option	Description
Open	Opens the Master File for viewing and editing using a graphical interface. If an Access file is used it will be also available.
Sample Data	Retrieves up to 50 rows from the associated data source.
Data Profiling	<p>Data Profiling provides the data characteristics for synonym columns. Reports are available for Statistics, Count, and Key Analysis.</p> <p>Alphanumeric columns provide the count of distinct values, total count, maximum, minimum, average length, and number of nulls.</p> <p>Numeric columns provide the count of distinct values, total count, maximum, minimum, average value, and number of nulls.</p>
Analysis Reports	<p>Provides the following options.</p> <p>Impact Analysis. Generates a report showing where this synonym is stored and used, with links to the synonym instances. Impact Analysis reports enable you to evaluate changes before they are made by showing which components will be affected. For information about impact analysis reports, see Running Impact Analysis Reports on page 569.</p> <p>Dependencies Analysis. Generates a report showing information about the synonym and other synonyms and objects that are referenced within it.</p>

Option	Description
<p>Metadata Management</p>	<p>Options are:</p> <p>Create Cluster Synonym. Allows you to create a cluster synonym.</p> <p>Refresh. Opens the Update Synonym page.</p> <p>Edit as Text. Enables you to view and manually edit the Master File synonym. To update the synonym, it is strongly recommended that you use the graphical interface provided by the <i>Open</i> option, rather than manually editing the Master File.</p> <p>Edit Access File as Text. Enables you to view and manually edit the Access File synonym. This option is available only when an Access File is created as part of the synonym.</p> <p>Download Master File. Downloads the Master File to your PC.</p> <p>Download Access File. Downloads the Access File to your PC.</p> <p>Prepare Translation Files. Prepares metadata translation files, as described in LNGPREP Utility: Preparing Metadata Language Files on page 252.</p>

Option	Description
Data Management	<p>Followed by these options, if applicable:</p> <p>Recreate DBMS Table. Recreates the data source table. You are asked to confirm this selection before the table is regenerated. (Note that the table will be dropped and recreated. During the process, data may be lost.)</p> <p>Delete All Data. Deletes all existing data. You are asked to confirm this selection before the data is deleted.</p> <p>Drop Table. Drops the table so that it is removed from the DBMS.</p> <p>Insert Sample Data. Inserts specified number of sample records, populating all fields with counter values. If the synonym is for an adapter that supports Bulk Load, a <i>Bulk Load</i> check box is available in the Insert Sample Data dialog. If you select the Bulk Load check box and click <i>Insert</i>, the sample data will be inserted using Bulk Load. This reduces load times when a large number of sample rows are inserted.</p> <p>Show/Modify Data. Opens a window that shows the data in the data source with buttons you can click to insert values, filter values, reload the data source, and customize the view.</p> <p>Reorganize. Recreates the data source table preserving original data. This option is not available in the Reporting Server browser interface.</p> <p>Quick Copy. Enables you to copy the data to a table in a target adapter.</p>
Flow	Opens the data flow page for creating flows.
Copy	Copies the synonym to the clipboard.
Delete	Deletes the synonym. You are asked to confirm this selection before the synonym is deleted.
Cut	Deletes the synonym and places it on the clipboard.
Rename	Renames the synonym.

Option	Description
Properties	<p>Displays the properties of the synonym, including physical location, last modified date, description, and privileges.</p> <p>Click <i>Check File</i> to check the validity of the synonym.</p>

LNGPREP Utility: Preparing Metadata Language Files

The LNGPREP utility extracts TITLE, DESCRIPTION, CAPTION, and PROMPT attribute values from application Master Files into specially formatted language translation files for each language you need. Once you have the contents of these language files translated, your users can run these applications in the language they choose.

LNGPREP does two things. It extracts attribute values from a Master File into language files, and it inserts or updates the TRANS_FILE attribute in the Master File with a value identifying the application folder where the language files reside and a prefix used for naming the set of language files. If the Master File is part of a cluster, LNGPREP will extract translatable strings from every Master File referenced in the cluster, and will update each with the same TRANS_FILE value.

LNGPREP requires an input file listing the three-character codes of the languages you need.

The name of each language file starts with the prefix specified in the TRANS_FILE value, followed by a three-character language code, and the extension *.lng*.

For example, assume the language input file contains the French and Spanish language codes:

```
fre
spa
```

If the Master File specifies:

```
trans_file = xlate/xl_
```

The language translation files would be in the *xlate* application folder, named:

```
xl_fre.lng
```

for French

```
xl_spa.lng
```

for Spanish

Reference: The Base Language File

Each Master File must have a single base language in which the DESCRIPTION, TITLE, CAPTION, and PROMPT attributes are specified. This language need not be English.

LNGPREP extracts these attribute values into the base language file, whose language code, for historical reasons, is *eng*. In this case, *eng* does *not* mean English. It means *whatever language the Master File is written in*.

The base language file (*prefixeng.lng*) should *never* be hand edited. All other lng files *must* be hand edited by translators, to translate the string values from the base language to the appropriate language.

Translating Applications into English

Since language code *eng* is reserved to mean base language, you cannot use it to contain English translations of an application whose base language is not English. In those cases, use any of the other English dialect language codes, such as AME, UKE, CAE, or AUE. For example, if the base language is German, specify AME in the languages file, run LNGPREP, and it will produce *prefixeng.lng* and *prefixame.lng* files, both in German. Translate the contents of *prefixame.lng* into English. Leave *prefixeng.lng* untouched.

Reference: How Translated Master File Attributes Are Displayed

Each language file contains a line for each attribute value from a related set of Master Files. Each attribute value has a unique index number assigned to it. For example, if the Master File contains FIELDNAME=PRODUCT_CATEGORY, TITLE='Product,Category', and that TITLE happens to be the 39th translatable attribute value, LNGPREP will produce lng files all containing the line:

```
39 = Product,Category
```

Your French translator will edit *prefixfre.lng*, leaving the index values unchanged while translating the string values, producing, in this case,

```
39 = Produit,Catégorie
```

At run time, when the TITLE for field PRODUCT_CATEGORY needs to be displayed, if WebFOCUS is configured for LANG=FRE, WebFOCUS looks up "Product,Category" in *prefixeng.lng*, finds index value 39, looks up 39 in *prefixfre.lng*, and displays the TITLE as "Produit,Catégorie."

LNGPREP Modes

You can run LNGPREP from the Reporting Server browser interface using the *Prepare Translation Files* option, or you can run it using syntax. In either case, you must first create a configuration file containing the three-character language codes for each translation file you need, one language code on each line. The first invocation of LNGPREP for a given Master File adds the TRANS_FILE attribute in that and all related Master Files, creates the base language file by scanning the Master Files for supported attribute values, and creates a copy of the base language file with the correct name for each additional language. Then, a translator has to translate the values in each additional language file from the base language to the correct language for that file.

On each subsequent run, LNGPREP will check for updates to the list of related Master Files and attribute values and update the files as needed. Translators will then have to translate any attribute values added to the language files.

LNGPREP Best Practice

The recommended best practice is to create an app directory solely for the purpose of containing .lng files, and use this appname and a common prefix value for all LNGPREP commands. In addition, put the languages *fn.cfg* file in this app folder. This will create one set of .lng files for all apps, minimizing the time and effort spent on translation.

Procedure: How to Prepare Metadata Language Files Using the Reporting Server Browser Interface

1. Right-click a synonym, point to *Metadata Management*, then click *Prepare Translation Files*.

The Set Translation Files page opens, as shown in the following image.

The image shows a dialog box titled 'Set Translation Files'. It has three input fields:

- 'Application for Translation Files': contains the text 'ibisamp' and an ellipsis button to its right.
- 'Prefix': is an empty text box.
- 'Languages File': is an empty text box with an ellipsis button to its right.

 At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

2. Enter the following values or accept the defaults.

Application for Translation Files

Is the name of the application where the language files will be stored. You can click the ellipsis to select an application from the current application path. By default, it is the application where the synonym resides.

Prefix

Is the prefix value for the translation files for the selected synonym.

Languages File

Is the file containing the list of language codes for which translation files should be prepared. The file must have the extension `.cfg`, be stored in an application directory on the application path, and have one language code on each line. You can click the ellipsis to select the application where the languages file is stored.

3. Click *OK*.

The language files are prepared using the application, prefix, and languages configuration file you specified. A status page will open listing the language files created and the Master Files processed.

Syntax: How to Run the LNGPREP Command Using Syntax

```
LNGPREP FILE n_part_name LNGAPP appname LNGPREFIX prefix
          LNGFILE appname/fn
```

where:

n_part_name

Specifies the n-part (app1/app2...) name of a Master File.

appname

Specifies the location where `.lng` files will be written and updated.

prefix

Specifies the literal characters that will precede the three-character language code in the names of the `.lng` files.

appname/fn

Specifies the *appname* and *filename* of a user-created `.cfg` file containing the list of three-character language codes, one per line. For example, the following file named `langretail.cfg` contains language codes for American English, French, and Japanese:

```
ame
fre
jpn
```

Example: Sample LNGPREP Command

Assume the `Inglst.cfg` file contains the language codes `fre` (French) and `spa` (Spanish):

```
fre
spa
```

Uploading Files

Issue the following LNGPREP command:

```
LNGPREP FILE weather/forecast LNGAPP xlate LNGPREFIX tq_ LNGFILE xlate/lnclist
```

Alternately, you can right-click the forecast synonym, point to *Metadata Management*, and click *Prepare Translation Files*. The *Set Translation File for weather/forecast.mas* page opens. Enter the values shown in the following image and click *OK*.



Application for Translation Files: xlate ...

Prefix: tq_

Languages File: xlate/lnclist.cfg ...

OK Cancel

The following language files will be created:

- xlate / tq_eng.lng
- xlate / tq_fre.lng
- xlate / tq_spa.lng

The Master File weather/forecast.mas will be updated with the following attribute

```
TRANS_FILE= xlate/tq_
```

Translators then have to translate the values in xlate/tq_fre.lng and xlate/tq_spa.lng.

Uploading Files

The Get Data page enables you to upload data files from your local machine to server application folders and to load them into a DBMS for use in creating synonyms and reports. You can then run a report with or without email, or schedule it.

You can configure the upload process to validate and upload a specific list of file extensions.

The upload process enables you to:

- Prepare files for upload (store them in a temporary location for editing prior to upload).
- Edit the synonyms generated for the prepared files.
- Return to the console.

Note: Upload requires nested application directories to be enabled. By default, files are uploaded to the myhome application for the connected user (homeapps/secprovider_userid). For information about nested application directories, see [How to Set the Level of Nested Application Directories](#) on page 151.

Reference: Upload Support for Validation of File Extensions

You can configure uploads to operate on a specific list of file extensions.

1. On the Applications page, click *Manage* on the ribbon and click *Settings* from the shortcut menu.

By default, the upload_allowed entry field contains a comma-separated list of all file extensions that the server can upload.

2. Edit the list to contain only the file extensions that should be allowed.
3. Click Save.

After the server has restarted, the Upload process will not show files with unsupported extensions.

Preparing Spreadsheets for Upload

Uploading data to WebFOCUS can be made easier if you first familiarize yourself with the data file, and then ensure that it is properly formatted for upload, so that geographic data, hierarchies, and other important aspects of your data are recognized. This is important because the synonym created for your uploaded data provides the basis for quality analytical content.

You can use the following techniques to prepare your data for the uploading process.

Naming Conventions and Microsoft Excel Sheet Names


For Excel spreadsheets, the name of the file is not important, but the name of the worksheet that contains your data is used to generate the synonym name. For the best results, note the following guidelines:

- ❑ Ensure that the worksheet name is meaningful. For example, Store Sales is better than Sheet1.
- ❑ Remove special characters from the worksheet name. Spaces will be converted to underscores, but all non-alphanumeric characters should be removed from the name, such as the following:

```
/^~!@#$$%^()+=}{};,
```

- Remove or replace NLS characters with standard alphanumeric characters.

The following image shows a worksheet with a meaningful name, Retail Sales.



While you have an opportunity to edit the worksheet name within the Upload workflow, editing in Excel may be preferable. Note that your spreadsheet data and column titles may contain National Language Support data and special characters.

For CSV files, there is no worksheet name, so the CSV file name is used to generate the synonym name. For this reason, all of the limitations identified for Excel worksheet names apply to the CSV file name. Be sure to check and adjust the file name prior to the upload.

Removing Introductory Information

Sometimes, an Excel spreadsheet contains formatted headings in the first few rows. This information cannot be imported into WebFOCUS and should be removed. Delete the introductory rows and save the file before uploading. Alternatively, you can define a data range within your worksheet and leave the introductory information in place. The following image shows an example spreadsheet with a heading and subheadings highlighted.

The image shows an Excel spreadsheet with the following data:

	A	B	C	D
1	U.S. Energy Information Administration			
2	<i>May 2014 Monthly Energy Review</i>			
3				
4	Release Date: May 28, 2014			
5	Next Update: June 25, 2014			
6				
7	Annual Total	Primary Energy Consumed by the Residential Sector (Trillion Btu)	Total Energy Consumed by the Residential Sector (Trillion Btu)	Primary Energy Consumed by the Commercial Sector (Trillion Btu)
8				
9	1949	4460.434	5599.25	2668.869
10	1950	4829.337	5988.553	2834.094
11	1951	5104.476	6380.193	2737.684
12	1952	5158.193	6560.243	2672.911

The introductory information (rows 1-6) is highlighted with a red dashed box.

Placing Column Titles in the First Row

For data to be useful in WebFOCUS, your data columns must be identified and properly described in the synonym that is generated during the upload process. You can make this easier by ensuring that the first few rows of your Excel spreadsheet contain column titles that are meaningful to you and to other users who will be using it. An example of meaningful column titles is shown in the following image.

	A	B	C	D	E	F	G	H
1	Store Name	Country	State	City	Sale Date	Revenue	Cost of Goods	Gross Profit
2	Des Moines	United States	Iowa	Des Moines	9/5/2013	3487.45	2288	1199.45
3	Dayton	United States	Ohio	Dayton	9/1/2013	13371.56	9459	3912.56
4	South Salt Lake	United States	Utah	South Salt Lake	9/2/2013	5642.79	4049	1593.79
5	Oslo	Norway	Oslo	Oslo	10/2/2013	13094.04	9206	3888.04
6	Des Moines	United States	Iowa	Des Moines	10/5/2013	4821.17	2962	1859.17
7	Dayton	United States	Ohio	Dayton	10/2/2013	13094.04	9206	3888.04
8	South Salt Lake	United States	Utah	South Salt Lake	10/5/2013	5564.52	3853	1711.52
9	Oslo	Norway	Oslo	Oslo	11/13/2013	18957.63	13293	5664.63
10	Des Moines	United States	Iowa	Des Moines	11/6/2013	6092.9	4513	1579.9

If your spreadsheet has more than one row of column titles, WebFOCUS can merge the information when creating the synonym. You will be given an option to specify how many first rows of the Excel file contain title information in the Upload workflow.

Removing Aggregated Information

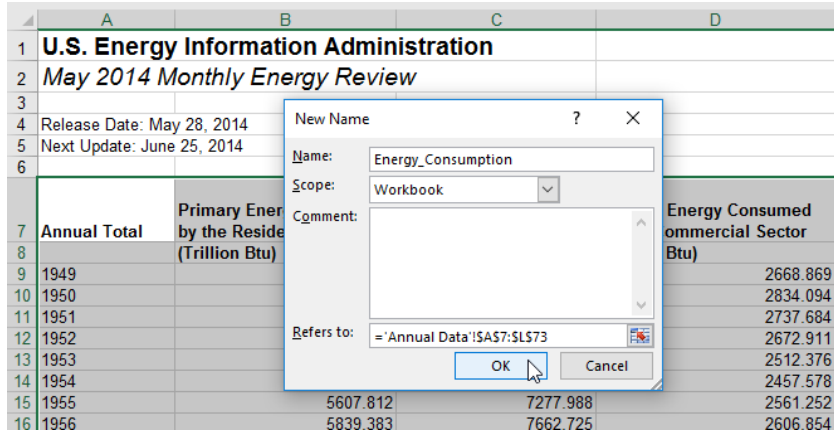
Excel spreadsheets may contain subtotals, grand totals, and other non-data row information. Data aggregation is performed by WebFOCUS, so you should remove these kinds of rows from your spreadsheet and save the file before uploading it.

Using Excel Name Ranges

Data ranges defined within your Excel worksheet can be helpful for the following reasons:

- Your worksheet may have introductory information, such as formatted headings or non-data information, in the first few columns.
- You may not want to import all of the data columns found on your worksheet.

You can define a data range in your worksheet to remove the data that you want WebFOCUS to process during upload, and leave your spreadsheet in its original format. An example of this is shown in the following image.



Preparing Hierarchical Data Columns

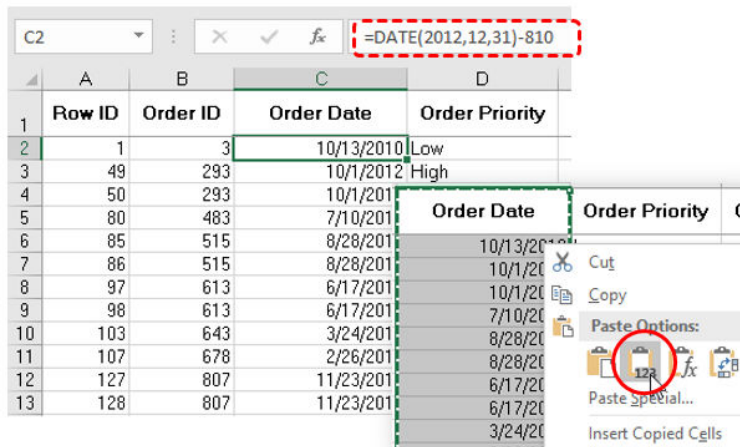
WebFOCUS recognizes columns of data that have hierarchical relationships. This is useful because the field names are arranged more logically in InfoAssist and because it facilitates Auto Drill capabilities in the content. Auto Drill allows you to drill up and down a field hierarchy automatically, making the content engaging and useful.

To help WebFOCUS recognize hierarchical columns correctly, ensure that the column titles begin with a common word or words, and are arranged left-to-right in the correct top-to-bottom direction, as shown in the following image. In the Upload workflow, you can define and edit dimension hierarchies prior to creating the synonym. You can also do this prior to the upload in Excel.

	A	B	C	D	E	F	G
1	Store Name	Product Category	Product Subcategory	Sale Date	Revenue	Cost of Goods	Gross Profit
2	Des Moines	Video Production	Video Editing	9/5/2013	3487.45	2288	1199.45
3	Dayton	Video Production	Video Editing	9/1/2013	13371.56	9459	3912.56
4	Dayton	Video Production	Video Editing	12/1/2013	19097.49	13386	5711.49
5	South Salt Lake	Video Production	Video Editing	12/1/2013	10763.83	8138	2625.83
6	Brussels	Accessories	Universal Remote Controls	3/2/2010	10060.72	7822	2238.72
7	Brussels	Accessories	Universal Remote Controls	4/1/2010	12839.15	9456	3383.15
8	Brussels	Accessories	Universal Remote Controls	5/1/2010	9848.13	6992	2856.13
9	Brussels	Accessories	Universal Remote Controls	6/1/2010	10357.7	7584	2773.7
10	Brussels	Accessories	Universal Remote Controls	7/1/2010	10979.11	8283	2696.11

Removing Date Formulas

Spreadsheets may contain a date column where the values are computed by Excel using a formula. You need to convert these computed values into simple values before uploading a spreadsheet. To remove date formulas, select a column, right-click, and then click *Copy*. Then, right-click the selected column again and select *Values*. This can be found under the Paste Options menu, which is highlighted in the following image. Menu options may vary by Excel release. Now you can see that each cell contains a date value, and WebFOCUS can decompose your dates into useful components for use in InfoAssist.



Starting the Upload Process

If you are about to upload an Excel file, make sure you have followed the information about preparing data for upload described in [Preparing Spreadsheets for Upload](#) on page 257.

Click a file type in the Desktop files section of the initial Get Data page or right-click a file type on the Advanced Mode page and click *Upload Data*.

A file picker dialog box opens.

Navigate to a file on your file system, and click *Open*.

if you started from the Advanced Mode page, you can enhance the synonym, as described in [Using Data Assist](#) on page 270.

On the Upload page, you can select the adapter and application in which you want to upload the file.

Click *Load and Report*.

If you started from the Simple Mode page, you can preview the data or go on to create a query. If you started from the Advanced Mode page, the Load Options dialog box opens so that you can select an adapter for loading the data and options specific to the adapter, as shown in the following image.

The screenshot shows the 'Load Options' dialog box. It features a title bar with the text 'Load Options' and a close button (X) on the right. The dialog contains several configuration options, each with a question mark icon to its left:

- Load Option:** A dropdown menu set to 'New/Replace'.
- Adapter:** A dropdown menu set to 'MS SQL Server ODBC'.
- Connection:** A dropdown menu set to 'MS SQL Server ODBC'.
- Synonym Application:** A dropdown menu set to 'Delimited Files (CSV/TAB)'. To its right is a text input field with an ellipsis button.
- Synonym:** A text input field set to 'DATREC - fast binary'.
- Table Name:** A text input field set to 'JSON STRUCTURED'.
- Bulk Load:** A dropdown menu set to 'XLSX - Excel'.
- Key columns derived from:** A dropdown menu set to 'No Key'.
- Apply Filters:** A checkbox that is checked.

At the bottom right of the dialog are two buttons: 'OK' (highlighted in blue) and 'Cancel'.

You can also set default values for these load options, as described in [Displaying Settings for Reporting Server Browser Interface Preferences](#) on page 601.

Note: If bulk load is not available for the target adapter, un-click the Bulk Load check box.

When you are finished selecting your load options, click *OK*.

To exit the Upload process at any point, click X at the top right of the page.

A message similar to the following is added to the edaprint log file for each upload:

```
I request by cmrph000063 to upload file=file.ext,size=size,u=security/userid
```

Classifying Metadata During Upload

Metadata Classification examines your data and assigns classifications to the columns, which can then be used to match columns from separate data sources.

When uploading a data file, you can view the recommended metadata classification values for each character-valued column. You can choose to keep the recommended values, or to change them.

Including classification values in your data improves the accuracy of correctly mapping table columns to each other. This is useful if you are integrating data from multiple sources, or if integrating data into a system with a predefined hierarchy.

The Data Classification engine is a machine learning application, provided as a service, that will base class assignments on examining actual data values. It is trained by being given access to large volumes of data gathered throughout the world. This application will continually learn more about data and how to classify it, and new classes will be added as needed.

Prerequisites for Metadata Classification

In prior releases, automatic matching of columns from multiple data sources was based either on primary key/foreign key relationships in relational tables or, if those were not available, column names and formats. The new Data Classification engine is a machine learning application, provided as a service, that will base class assignments on examining actual data values. It is trained by being given access to large volumes of data gathered throughout the world. This application will continually learn more about data and how to classify it, and new classes will be added as needed.

A customer site wanting to use the Data Classification application will have to submit proprietary data to the classification API. To ensure the safety and integrity of proprietary data, the API will be installed as a service in the customer environment, and will be accessible only to the customer. Once the data is examined, feature vectors with encoded values will be sent to another service, available in the WebFOCUS or TIBCO Omni-Gen® environment, which will use these vectors to recommend class assignments for columns. If a classification cannot be recommended based on the known classes, an attempt will be made to classify columns based on column names and formats. You will be able to view and change the classifications prior to saving them in a synonym. Once they are stored in a synonym, they will be used to suggest column matches when you create a Union in a Data Flow on the server. In future releases, the functionality will be extended to other environments.

Instructions for installing the services are included in your *TIBCO WebFOCUS® Installation and Configuration* manual. Please make sure your installation team has installed them prior to using this feature.

After the services are installed, you must add the paths to the services in the `edaserve.cfg` file on the server.

Changes Needed in the edaserv.cfg Configuration File

The URLs pointing to the services (APIs) must be added to the edaserve.cfg configuration file on the WebFOCUS Reporting Server.

1. On the Reporting Server browser interface, navigate to the Workspace page using the Tools menu.
2. On the Resources Tree, expand *Configuration Files*.
3. Right-click *Workspace - edaserve.cfg* and click *Edit*.

The edaserve.cfg file opens in the Text Editor.

4. Enter values for the following keywords:

```
lblfeat_url = url_to_service_that_encodes_values
lblpred_url = url_to_service_calculates_labels_based_on_encoded_values
```

where:

url_to_service_that_encodes_values

Is the URL to the service that resides in the customer environment. For example:

```
lblfeat_url = https://myserver/api/1.0.0
```

This service is called by passing values of rows in a column, and it passes back a feature vector that describes the patterns of values in the column.


url_to_service_calculates_labels_based_on_encoded_values

Is the URL to the service that resides in the WebFOCUS and Omni-Gen® environment. For example:

```
lblpred_url = https://ibiserver/api/1.0.0
```

This service is called using the feature vector created by the first service and returns the classification for the column.

The Text Editor should look similar to the following.



- Click the Save icon and close the tab.

Procedure: How to Change Metadata Classification Values When Uploading Data Files

You can change the metadata classification values for your data during the upload process.

- From the WebFOCUS start page or Server Get Data page, upload an Excel or delimited file using Get Data.

A preview page opens, as shown in the following image.

The screenshot shows the 'Uploading Data' window with the 'SALES_ORDERS (PREVIEW/MODIFY TABLE)' section. The table below is a preview of the data:

Order Date	Sales Rep	Store Code	Store Name	Country	Region	State	City	Product Number	Product Type	Product Category	Model
10/17/2017	Web	9999CE	Century Electronics (Web)	United States	Web(R)	Web(S)	Web(B)	2005	Video	DVD	DVD-150P
10/17/2017	Web	9999CE	Century Electronics (Web)	United States	Web(R)	Web(S)	Web(C)	3004	Cameras	Digital Cameras	DC-610
10/17/2017	Web	9999CE	Century Electronics (Web)	United States	Web(R)	Web(S)	Web(C)	4022	Camcorders	Digital Camcorders	CDV-1600
10/17/2017	Web	9999CE	Century Electronics (Web)	United States	Web(R)	Web(S)	Web(C)	5002	Office	Organizers	EO-100

- Click the *Classify* button.

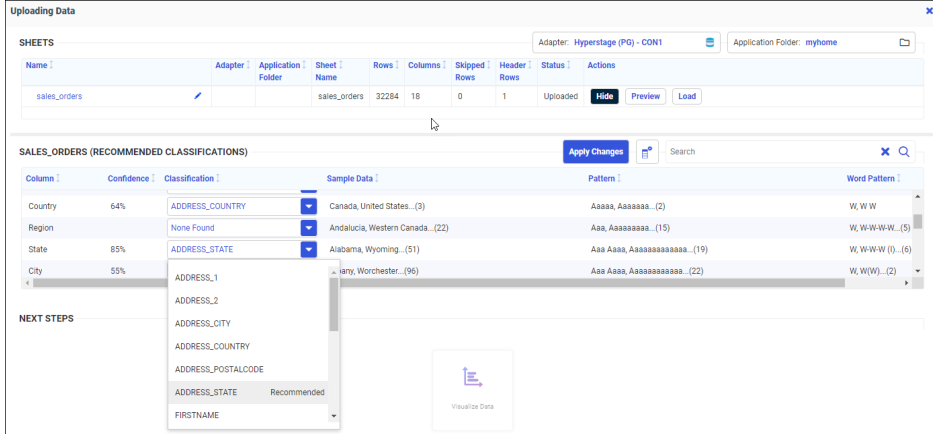
The Recommended Classifications page opens, as shown in the following image.

The screenshot shows the 'Uploading Data' window with the 'SALES_ORDERS (RECOMMENDED CLASSIFICATIONS)' section. The table below shows the recommended classifications:

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
Country	64%	ADDRESS_COUNTRY	Canada, United States... (3)	Aaaaa, Aa33333... (2)	W, W, W
Region	None Found	None Found	Andalusia, Western Canada... (22)	Aaa, Aaaa33333... (15)	W, W-W-W-W... (5)
State	85%	ADDRESS_STATE	Alabama, Wyoming... (51)	Aaa Aaaa, Aa3333333333... (19)	W, W-W-W... (6)
City	55%	ADDRESS_CITY	Albany, Worcester... (96)	Aaa Aaaa, Aa3333333333... (22)	W, W(W)... (2)


When a recommended classification can be made, the columns that were classified have a Confidence Level associated with them, which indicates the probability that the classification is correct.

You can select from a column list of classifications, or enter your own classification, by typing into the text box next to the column. Clicking a column drop-down arrow shows the list of classifications, and indicates which is recommended, as shown in the following image.



3. Click *Apply Changes*.

Your changes are applied, as indicated by a green message bar at the top of the page, as shown in the following image.



Once your data is loaded, with the proper permissions to edit synonyms, you can leverage further capabilities by accessing the server.

Getting Data

The Get Data page provides a process for connecting to data and creating a base synonym or cluster synonym. It can be used to create metadata for different types of data sources. Some of the options and parameters vary depending on the type of data source. The Connect to Data page enables you to upload data (as described in [Uploading Files](#) on page 256) or configure an adapter, create a new connection for an adapter, or change adapter connection parameters.

To return to the full console, click X at the top right of the page.

Procedure: How to Use the Get Data Page to Configure an Adapter, Create a New Connection, or Change a Connection

In the Server Files section of the Get Data page, you can perform one of the following tasks.

1. To configure a new adapter, click the New Datasource button (+).

A page of available adapters opens. Enter the connection parameters.

2. Right-click an adapter and click *Configure*

The Connection Parameters page for the adapter opens.

3. Enter the connection parameters.

Note: Connection parameters for each adapter are described in the *TIBCO WebFOCUS® Adapter Administration* manual.

4. To add a connection for a configured adapter, right-click an adapter on the Advanced Mode page and click *Add Connection*.

The Add Connection page opens. Enter the connection parameters.

5. To change connection parameters, right-click a connection on the Server Files list and click *Properties*.

Note: To show the configured connections for an adapter on the Simple Mode page right-click the adapter and click *Show Connections*.

The Change Connect Parameters page opens. Change the connection parameters.

6. To duplicate a connection, right-click a connection and click *Duplicate Connection*.

The Duplicate Connection page opens with parameters entered for the same adapter and the same server as the original connection. You can use this to connect to a different database than the one in the original connection, or use a different type of security model.

Procedure: How to Import Metadata

After completing the Connection step, you can create synonyms for configured adapters and connections.

1. To create synonyms, right-click a connection and click *Show local files* (for delimited files) or *Show DBMS objects* from the context menu.

For an SQL adapter, the Available Objects page opens. If you started from Advanced Mode, you have three options.

- Cluster Synonym.** All selected tables are added, and join columns can be selected as needed. This is the default selection.
- Cluster Join for Star Schema.** Only creates join conditions when the selected tables have Primary Key/Foreign Key relationships.
- Base Synonym.** Creates a synonym for an individual table.

The following image shows *Cluster Join for Star Schema*, with one fact table and two dimension tables selected.

Create Synonym for MS SQL Server ODBC/AzureDB (wfiretail)

Object Type? Database? Owner/Schema? Object Name?

Customize data type mappings

Create? Cluster Synonym Cluster Join for Star Schema Base Synonym

Application? Prefix? Suffix?

SYNONYM CANDIDATES FOR DATABASE: Wfiretail (DEFAULT) Row Limit: Search:

Fact	Dimension	Default Synonym Name	Table Name	Owner/Schema	Type
<input checked="" type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_sales_s	wrd_wf_retail_sales	dbo	TABLE
<input type="checkbox"/>	<input checked="" type="checkbox"/>	wrd_wf_retail_product_s	wrd_wf_retail_product	dbo	TABLE
<input type="checkbox"/>	<input checked="" type="checkbox"/>	wrd_wf_retail_store_s	wrd_wf_retail_store	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_age	wrd_wf_retail_age	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_currency	wrd_wf_retail_currency	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_customer	wrd_wf_retail_customer	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_customer_missingdata	wrd_wf_retail_customer_missingdata	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_customer_ss1	wrd_wf_retail_customer_ss1	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_customer_ss2	wrd_wf_retail_customer_ss2	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_discount	wrd_wf_retail_discount	dbo	TABLE
<input type="checkbox"/>	<input type="checkbox"/>	wrd_wf_retail_education	wrd_wf_retail_education	dbo	TABLE

The Show DBMS Objects page has been designed so that you can enter all the parameters needed to create base synonyms or a cluster synonym on a single page.

2. For an SQL adapter, the Filter section enables you to select the object type and enter the owner/schema and database from which to generate the list of objects for selection.

You can use the object name box and Search button to list specific object names.

- To create a base synonym, click the *Base Synonym* radio button and click the check box next to the object for which you want to create a synonym.

When you are finished, click *Add*.

- To create a cluster synonym, click the *Cluster Synonym* or *Cluster Join for Star Schema* radio button and click the check box next to the fact and dimension tables to be part of the cluster.

The DBMS objects list enables you to select objects to add to the cluster and assign them as facts or dimensions.

You can select each fact and each dimension manually, or you can right-click a fact or dimension table name and select one of the following options, which are available if the foreign key to primary key information is available in the RDBMS:

- Show Related Dimensions.** Displays a report of related dimensions, as shown in the following image.

The screenshot shows the 'Available Objects for MS SQL Server ODBC (CON01)' dialog box. The 'Create' section has 'Cluster Synonym' selected. The 'Table Name' list shows 'wrд_wf_retail_sales' selected with a checkmark. A side panel titled 'Dimensions related to wrд_wf_retail_sales' displays a table of related dimensions.

PK_TABLE_NAME	FK_NAME
wrd_wf_retail_currency	wrd_fact_sales_id_currency_fk
wrd_wf_retail_customer	wrd_fact_sales_id_customer_fk
wrd_wf_retail_discount	wrd_fact_sales_id_discount_fk
wrd_wf_retail_product	wrd_fact_sales_id_product_fk
wrd_wf_retail_store	wrd_fact_sales_id_store_fk
wrd_wf_retail_time	wrd_fact_sales_id_time_fk

- ❑ **Add Related Dimensions.** Selects all related dimensions automatically.

When you are finished, click *Add*.

Note: If you selected the *Cluster Join for Star Schema* radio button, and no primary key/foreign key relationship is found, the corresponding dimension table is not added to the cluster, and a diagnostic message is issued.

If you started from Simple Mode, the synonym is added in the application directory specified. If you started from Advanced Mode, Data Assist opens so that you can enhance the synonym and, optionally, specify load options, as described in [Using Data Assist](#) on page 270.

After the synonym is saved, you can edit the synonym by right-clicking the synonym on the application page and clicking *Open*. For detailed information about enhancing the synonym, see [Using Data Assist](#) on page 270.

Using Data Assist

Data Assist provides access to the Synonym Editor, which is an environment for managing and editing synonyms.

Introduction to Metadata

Metadata is data about data. Metadata describes the characteristics of your data sources and provides the information needed to access those data sources. This insulates you from the need to know how to optimally access the data, so that you can concentrate on using the data for reporting, charting, visualizations, and applications. Examples of metadata components are data type, data location, and data content and information about how to connect to it. Metadata can be generated whenever data is created, acquired, added to, deleted from, or updated in any data source within an enterprise.

Understanding Data Terminology

The smallest meaningful element of data in a data source is referred to as a column (for relational data sources) or a field (for non-relational data sources). Column and field are used interchangeably throughout this content. Every column or field has several characteristics, such as the data type (character, date, integer, or some other type) and length or scale. Field characteristics can also include output display options, such as comma inclusion, currency symbol, date display, and so on. For consistency, all fields should be included in the metadata, so all users get the same view of the data. However, you can create a logical view of the data in which only a subset of the fields are available.

Certain fields in a data source may have a one-to-one relationship with each other and describe a group of related characteristics. For example, each unique ID number represents only one employee. These fields can be grouped together into a segment or, in relational terms, a table. You can relate segments to each other by creating a multi-segment data source or by joining segments together. Segments are the building blocks of larger data structures.

Fields can be categorized as measures (facts), dimensions, hierarchy levels, or attributes.

A measure or fact is a numeric value, such as Gross Profit or Cost of Goods Sold, that you may want to aggregate. All numeric values that can logically be summed can be categorized as measures. Numeric fields that cannot be summed, such as product number and order ID, are not categorized as measures. Instead, they may be used in the same way as dimension fields to analyze measures. It is up to you to understand your data and determine whether each numeric field should be summed.

A dimension is a way to categorize data or sort output. You can use a dimension to analyze and compare measures. Generally, dimension fields are alphanumeric fields such as product. Certain dimension fields can be organized into hierarchies to define the relationships between the fields in the hierarchies and provide the information needed for the AUTODRILL feature. For example, a Geography hierarchy can contain the Continent, Country, State, and City fields. On report output, with AUTODRILL turned on, links will be provided on report or chart output that let you drill down the hierarchy levels to more detailed reports or charts. You can also define dimension fields that are not part of a dimension hierarchy.

An attribute is a field or collection of fields that add additional information about a dimension.

Using the Business View feature, you can create a view of the data that represents its business use by creating virtual segments called folders and adding fields to the folders. You can reuse fields in multiple folders. You can then categorize those folders and fields as measures (facts), dimensions, hierarchy levels, or attributes. Only the folder structure, not the real segments will be visible in the WebFOCUS tools and be available for use in requests.

Accessing Data

When WebFOCUS accesses data, it needs to know how to interpret the data that it finds. A Master File, which is the primary component of metadata, is a map of the segments in the data source and all of the fields in each segment. The Master File also identifies the name of the data source and the type of data source it is.

For some types of data sources, the Master File is supplemented by an Access File. An Access File includes additional information that completes the description of the data source. For example, it includes the full data source name and location. You require one Master File and, for some data sources, one Access File to describe a data source. These metadata files (Master File and Access File) are referred to as a synonym.

Synonyms are stored separately, apart from the associated data source. WebFOCUS uses a synonym to interpret a data source as follows:

1. It identifies the data sources available and their respective fields.
2. It examines the security rules, if internal DBA data source security has been specified for the data source, and ensures that user access is allowed based on any DBA security specified.
3. It identifies the Access File for the data source, if that data source requires an Access File.

The data source content is interpreted based on the information in the Master File and, if applicable, the Access File.

Working With Data Adapters

With the appropriate WebFOCUS Data Adapters configured, you can access data from many different sources, including the following:

- Relational data sources, such as Db2, Teradata, Oracle, and Sybase.
- Hierarchical data sources, such as IMS and FOCUS.
- Indexed data sources, such as ISAM and VSAM.
- Network data sources, such as CA-IDMS.
- Sequential data sources, both fixed-format and delimited format.
- Multi-dimensional data sources, such as SAP BW and Essbase.
- XML data sources.

Every adapter is specifically designed for the data source that it accesses, and, as a result, is able to translate between SQL or WebFOCUS and the Data Management Language of the data source. The adapter manages the communication between the data interface and the data source and returns either answer sets or messages to the requestor. Adapters provide solutions to product variations, including product differences in syntax, functionality, schema, data types, catalogs, data representations, message processing, and answer set retrieval. It is the adapter that manages the synonym creation process.

The parameters required to configure supported adapters and to create synonyms for the corresponding data sources are described in detail in the *TIBCO WebFOCUS® Adapter Administration* manual.

Generating Synonyms

Synonyms are useful because they insulate client applications from changes to the location and identity of a request. You can move or rename a request without modifying the client applications that use it. You need to make only one change, redefining the request synonym on the server. Synonyms provide support for the extended metadata features of the server, such as virtual fields and security mechanisms. Creating a synonym generates a Master File and an Access File.

Often, a synonym already exists on the server, and reporting can begin at once. However, if the synonym you require does not exist on the server, an authorized server or application administrator can create it directly using the Reporting Server browser interface, the Data Migrator desktop interface, or the Metadata canvas in TIBCO WebFOCUS® App Studio.

Although you can use Data Assist to create a new synonym, the situations in which you need to do this are extremely rare. The Reporting Server browser interface, Data Migrator desktop interface, and Metadata Canvas provide automatic and error-free synonym creation tools.

Once you have generated a synonym, you can report against the synonym. In many instances, the configured adapter and the generated synonym are all you need to access your data and create reports and graphs.

However, you may wish to enhance the synonym in order to implement particular capabilities that are supported in the WebFOCUS data description language. To do this, you can use Data Assist or the Metadata canvas.

Enhancing a Synonym

The following are some of the attributes you might want to add to the synonym to enhance your data access and reporting capabilities. You can:

- Add virtual columns (DEFINE fields) and columns for aggregated values (COMPUTE fields).
- Add filters to specify data selection criteria.
- Add group definitions for data sources that support groups.
- Add meaningful titles and descriptions, including multilanguage variations.
- Change the format of fields (for example, the size of an alphanumeric field or the format of a date field).

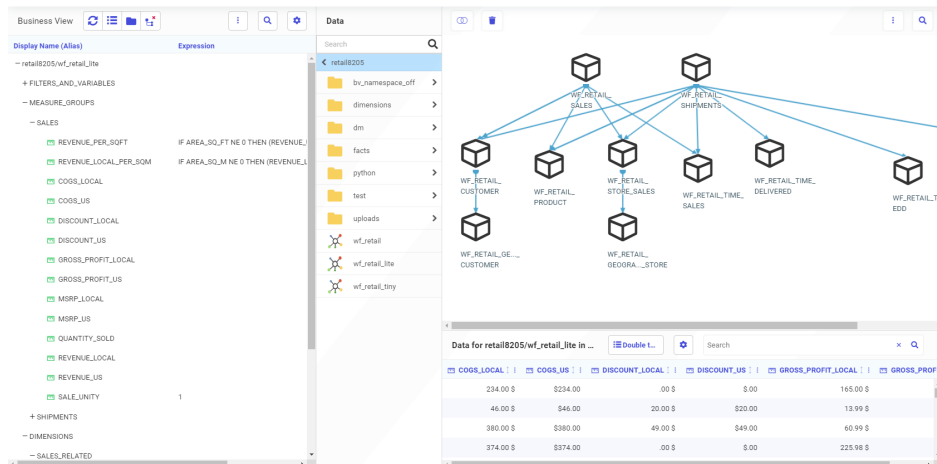
- Create a cluster join view by linking available synonyms to create a multisegment (multitable) file for reporting.
- Create business views of the metadata in order to limit the fields available to any retrieval request that references the business view and to group fields together based on their roles in an application.
- Define parent and child hierarchies for cube data sources.
- Apply security rules for fields and values to ensure that user access is based on internal data source security (DBA) specifications.

Opening the Synonym Editor

You open a synonym in the Synonym Editor by:

- Clicking *New* on the Applications page or right-clicking an application folder, clicking *New*, then *Cluster Business View*.
- Double-clicking a synonym on the Applications page.
- Right-clicking a synonym on the Applications page and clicking *Open*.
- Uploading a file or right-clicking an adapter connection on the Get Data page.

The synonym opens in the Synonym Editor, as shown in the following image of a Cluster Business View open in the Synonym Editor.



The Join Editor is open showing the joins included in the cluster, with an Applications pane you can use to add tables to the cluster.

The Output pane is open to show sample data.

Note that each data item has an icon that indicates the data type of that item. For information about:

- ❑ The Synonym Editor ribbon and frames, see [Understanding the Synonym Editor User Interface](#) on page 275.
- ❑ The enhancements you can make to a synonym, see [Enhancing a Synonym](#) on page 296.

Understanding the Synonym Editor User Interface

The Synonym Editor user interface consists of a ribbon and frames used for viewing and enhancing a synonym. You can open multiple synonyms in the Synonym Editor.

The number of frames you actually see when you open a synonym depends on the environment from which you are accessing the Synonym Editor and the characteristics of the synonym you open.

Reference: Synonym Editor Ribbon

The Synonym Editor Ribbon is shown in the following image. The buttons that display depend on the environment from which you opened the synonym.



Save

Saves the synonym.

Undo

Undoes the most recent actions. You can set the number of undo actions supported using the Options menu. The default number of undo actions is 50.

Redo

Redoes the most recent actions. You can set the number of redo actions supported using the Options menu. The default number of redo actions is 50.

View Menu

Reset View

Resets to the default view.

Show Table/Column

Opens the Table/Column pane, which shows the physical view of the segments and fields. This pane is hidden, by default.

Hide Business View

Closes the Business View pane, which shows the logical view of the synonym.

Hide Data

Hides the data pane.

Hide Join Editor

Hides the Join Editor frame.

Advanced Options Menu

Provides the following options.

Options

Opens the Advanced Options dialog box for setting the following preferences.

Synonym Editor

Includes the following options.

- Support extended options.** Selections are Yes and No. The default value is No. This will add the right-click options *sort objects* and *styles* to the appropriate fields and folders.

Column Management

Includes the following options.

- Name display strategy.** Allows you to select the type of metadata to display in the tree. Supported values are Name, Title, Description, Alias. The default value is Title.
- Function display strategy.** Allows you to select whether to display function syntax or a short description when you right-click a field and add a function. The default is a short description.

Language Generation

Includes the following options.

- Use segment to qualify field reference.** Defines how field references will be generated in requests. Selections are *For duplicate fields* and *Always*. The default value is *For duplicate fields*.
- Use application name to qualify synonym reference.** Defines how synonym references will be generated in requests. Selections are Yes and No. The default value is No.

Limits

Includes the following options.

- Maximum number of rows for test.** Enter the maximum number of rows for test queries. The default value is 50.
- Maximum number of columns for test.** Enter the maximum number of columns for test queries. The default value is 999999.
- Limit number of database reads for a filter.** Enter a limit for the number of initial database reads for a filter if performance for filters is an issue. You will then be able to edit the filter or add additional values.
- Maximum number of identical error messages.** Enter the maximum number of identical error messages that can be issued. The default value is 10.
- Random Sampling limit (in %).** Enter the maximum sample size for data profiling, as a percent.
- Undo/Redo Limit.** Enter the maximum number of undo/redo actions supported. The default value is 50.
- Omit missing (null) values in Data Profiling charts.** Select Yes to omit missing values in the charts generates for data profiling. The default value is No, which includes missing values.

DBA

Opens the DBA dialog box for adding security rules to the synonym. For more information, see [Adding Data Source Security \(DBA\)](#) on page 314.

AutoRun

Provides the following options.

Off

Does not show sample output.

On

Creates a tab in the output window with a sample of the current data. On is the default value.

Data Preview

Can be Off or On. On shows sample values in the Table/Columns and Business View panes. Off does not show sample data values in the Tables/Columns and Business View panes. On is the default value.

Sample Data

Allows you to generate a sample data report of the current synonym in the output pane.

Diagnostics

Provides the following options.

View MFD

Opens a text view of the Master File in the Output pane.

View AFD

Opens a text view of the Access File in the Output pane.

MFD Info

Displays a report about the segments and fields in the Master File.

Show Layout

Opens the layout window.

Clear Undo History

Clears the history of undo actions.

Help Menu

The Help menu (? icon) on the right of the ribbon provides the following options.

Contents and Search

Opens the Reporting Server browser interface help page.

Tutorials

Provides a tutorial for using Upload/Connect to Data and Business View +.

Version

Displays version information.

Browser Support

Opens a window that shows supported browsers and browser versions.

New Features

Opens a list of new feature manuals in PDF format for the various server releases.

Release Notes

Opens the Release Notes report for this version of the server.

Quick Setup Guides

Opens installation and adapter prerequisites help files.

Licenses

Enables you to open the TIBCO End User License Agreement or third-party licenses, if any.

Understanding Synonym Editor Frames

The Synonym Editor frames that you see by default depend on the environment from which you are opening a synonym and the characteristics of the synonym. The frames available in Synonym Editor are:

- ❑ **Table/Column frame.** This frame shows the physical structure (segments and fields) of the data source. It is hidden by default, but you can show it by clicking the *View* button, then *Show "Table/Column"*.
- ❑ **Business View frame.** This pane shows the logical view of the data source.
- ❑ **Filters/Groups frame.** This frame shows any filters and groups defined in the synonym or view. It only shows when you create a filter or group.
- ❑ **Output frame.** This frame shows sample data and reports.
- ❑ **Join frame.** This frame shows the joins defined in the cluster synonym in the modeling view. It also shows an Applications frame that you can use to add tables to the joins.

Each frame has a menu bar that offers options specific to that frame as well as options that are common to multiple frames.

Reference: Common Synonym Editor Options

The following menu bar options are common to multiple frames.

- ❑ The *Find* (magnifying glass) button is common to all multiple Synonym Editor frames. It toggles a text box for entering a search term.
After implementing a search, a cancel (undo) button becomes available to delete the search term.
- ❑ The Hover menu is common to all panes in the Reporting Server browser interface. You can open it by hovering the mouse on the top right of a pane. It enables you to drag the pane to a different location on the page, maximize the pane, close the pane, and specify layout options for the pane. The hover menu is shown above the Search icon on the Business View frame in the following image.



In addition, there are common right-click menu options for fields in the Table/Column, Business View, and Output panes. Some of the options that show up are dependent on the data type of the field. For more information, see [Right-Click Options for Fields in the Table/Column, Business View, and Output Panes](#) on page 321.

Reference: Table/Column Frame

The Table/Column frame displays the physical characteristics of the underlying data sources represented in the synonym. You can use options on the menu bar, or use the context-sensitive right-click menu options, to make changes to the physical synonym.

The Table/Column frame is shown in the following image.

Display Name (Title)	Expression
- retail8203/wf_retail_lite	
- Sales Fact	
ID Sales	
ID Store	
ID Currency	
ID Customer	
ID Discount	
ID Product	
ID Time	
Cost of Goods,Local Currency	
Cost of Goods	
Discount,Local Currency	
Discount	
Gross Profit,Local Currency	
Gross Profit	
MSRP,Local Currency	
MSRP	
Quantity,Sold	
Revenue,Local Currency	
Revenue	
fx Sale Unit(s)	1
fx Sale Unit(s)	1

The icon next to each column identifies the data type of the column.

The menu bar at the top of the pane has the following options.

- ❑ The *Display* menu (...) at the top of the pane gives you the following display options for the synonym tree.

[Columns/Tables](#)

Lists the segments and fields in the synonym with their data types and two data values. This is the default display.

[Tables](#)

Displays the segments (tables) and their relationships in the form of a tree.

[Join Editor](#)

Opens a frame that displays the Join Editor. You can edit the joins using menus and options at the top of the pane or right-click options. For more information, see [Creating a Join](#) on page 308.

- ❑ The *View* menu (down arrow) gives you the following options.

[Choose Columns ...](#)

Opens a dialog box for selecting the metadata attributes and number of field values to display.

[Reset to Defaults](#)

Resets the pane to the default view, showing the default columns.

Options When Right-Clicking the Synonym Name

Properties. Opens a dialog box for editing the following properties. Click *Apply* when you have finished editing the properties.

General Properties

The following general properties are available.

SUFFIX

identifies the type of data source being accessed. This cannot be changed in this dialog box.

MDF_PROFILE

Specifies a FOCEXEC to be run prior to any request that references the synonym.

FDEFCENT

Defines the default century for handling cross-century dates. Cannot be changed in this dialog box.

FYRTHRESH

Defines the year threshold for handling cross-century dates. Cannot be changed in this dialog box.

REMARKS

Provides a description.

Miscellaneous Properties

The following miscellaneous properties are available.

BV_NAMESPACE

Can be one of the following values. Switching the value can invalidate field references in a request.

- OFF.** Uses segment names to resolve field references. This is the default value.
- ON.** Is only available for upward compatibility with synonyms from prior releases. It should not be used except in consultation with Customer Support Services.

DV_INCLUDE

Can be one of the following values.

- YES.** Inherit dimension view properties from base synonyms in the cluster. This is the default value.
- NO.** Do not inherit dimension view properties from base synonyms in the cluster.

New Variable. Opens a dialog box for referencing a global or system variable in the synonym. Click *Apply* after entering or selecting values for the following properties.

Variable Name

Enter the name of a global or system variable to be referenced in the synonym.

Prompt Caption

Is descriptive text to replace the variable name in the prompt.

Default Value

Enter a default value for the variable.

Prompt Values

Select one of the following options.

- Input box.** The user will enter a value in a text box.
- Static list of values.** You define the list of values that the user can select from. Two additional fields open:
 - Select Type.** Select *Single selection*, if the user can only select one value from the list, or *Multiple selection*, if the user can multi-select values from the list.
 - Values.** Enter the list of values, enclosed in single quotation marks, and separated by the OR operator.
- Static list of values/captions.** You define the list of values that the user can select from and descriptive text to display instead of the values. Two additional fields open:
 - Select Type.** Select *Single selection*, if the user can only select one value from the list, or *Multiple selection*, if the user can multi-select values from the list.
 - Caption/Value Pairs.** Enter the list of captions and values, in the following form.


```
(<'cap1', 'val1'>,<'cap2', 'val2'>)
```
- Synonym based dynamic list of values.** The user will select from a dynamic list of values found in a field in a server-accessible data source. Four additional fields open:
 - Select Type.** Select *Single selection*, if the user can only select one value from the list, or *Multiple selection*, if the user can multi-select values from the list.
 - Lookup Synonym.** Enter the name of the synonym for the data source that contains the field values that will comprise the generated list of values, or click the ellipsis (...) to navigate to the synonym.
 - Lookup Field.** Select the name of the field that contains the acceptable values for the generated list.

Lookup Display Field. Select the name of the field that contains the values to be displayed on the generated list.

- Procedure based dynamic list of values.** The user will select from a dynamic list of values found in a WebFOCUS request. Five additional fields open:

Select Type. Select *Single selection*, if the user can only select one value from the list, or *Multiple selection*, if the user can multi-select values from the list.

Lookup Procedure. Enter the name of the procedure that contains the field values that will comprise the generated list of values, or click the ellipsis (...) to navigate to the procedure.

Lookup Procedure Parameters. Enter any parameters needed for running the procedure.

Lookup Field. Select the name of the field that contains the acceptable values for the generated list.

Lookup Display Field. Select the name of the field that contains the values to be displayed on the generated list.

Quoted

Select this option to enclose the string in single quotation marks.

Persistent

Can be one of the following values. ON is the default value for this option.

- ON.** The default value assigned to this Master File global amper variable cannot be changed, preventing the use of this variable as a filter selection in WebFOCUS Designer.
- OFF.** The default value assigned to this Master File global amper variable can be changed, permitting the use of this variable as a filter selection in WebFOCUS Designer.

Usage Format

Select from Integer, Character (fixed), or Decimal. Additional fields open to specify the length of the format and, for Decimal, the number of decimal places.

New Expression. Click *Advanced Define* to create a virtual field or *Advanced Compute* to create a calculated value. For information about creating expressions, see [Using the Expression Calculator](#) on page 301.

Data Profiling. Click *Statistics*, *Count*, or *Key Analysis*. For more information about data profiling reports, see [Other Synonym Enhancements](#) on page 321.

Sample Data. Click this option to generate sample data with the current version of the synonym.

Impact Analysis. Click this option to generate a report on where this synonym is used.


Reference: Business View Frame


The Business View (BV) frame shows the logical view of the synonym. You can drag fields from the Table/Column frame, use options on the menu bar, or use the context-sensitive right-click menu options to make changes to the nodes or folders.

If you open an existing synonym, and it has a BV, that BV will be shown. If the synonym does not have an existing BV, the BV is not built, and just the top node is shown.

Note: A setting named NTM_METADATA_FIELDS in the Data Assist section of the *Settings for Reporting Server browser interface Preferences* option on the Workspace page allows you to specify whether to show all fields or only those in the Business View folders of the synonym. By default, only fields in the Business View folders are visible.


If you do not open an existing synonym, the icons at the top of the BV frame control what is generated:

Create Default () creates a fully populated structure.

Create Template Folders () creates a DV structure that is not populated

If an existing synonym from a prior release has a BV structure, for upward compatibility it is still supported. However, it limits you to working with folders in the BV frame and requires procedures that use this synonym to use folder.field references instead of segment.field references.

The Business View frame is shown in the following image.



Note: if you need to drag a field to a location below the visible area of the screen, the frame will automatically scroll as you drag, so you can drop the field where you need to place it.

The menu bar at the top of the frame provides the following options. Note that certain options only appear when the view supports them.

- The *Reset Default* button resets the frame to the default view.
- The *Flatten* button removes child folders.
- The *Create Template Folders* button creates the pre-defined DV folders Filters, Measures, and Dimensions.

- ❑ The *Hide All* button removes the Business View. When the Business View is removed, a *Create Default* button appears to recreate the default Business View.
- ❑ The *Display* menu provides the following options for displaying the folder tree.

[Folders/Columns](#)

Lists the folders and fields in the view with their data types and two data values. This is the default display.

[Folders](#)

Displays the folders and their relationships in the form of a tree.

[Modeling View of Folders](#)

Displays the modeling view of the synonym, a pictorial view of the filters, join, and variables.

- ❑ The *View* menu (down arrow) gives you the following options.

[Choose Columns ...](#)

Opens a dialog box for selecting the metadata attributes and number of field values to display.

[Reset to Defaults](#)

Resets the frame to the default view.

If you right-click the cluster synonym name and click *Manage Folders*, you can create a new folder.


Reference: Filters/Groups Frame

The Filters/Groups frame shows filters and groups defined in the synonym when you create a new filter or edit an existing one.

Note: If you create a new Advanced Filter, the expression calculator opens for creating the expression that defines the filter. This is needed if the filter is complex, for example, if it is based on multiple fields. When you edit a complex filter, select *Advanced* from the *Edit Filter* menu. If you open a complex filter in the simple filter interface, it opens as read-only.

Each filter or group added is displayed in its own card (vertically tiled frame) within the Filters/Groups frame.

The Filters/Groups frame is shown in the following image.



Each group card has an expand button on the menu bar, which expands the Filters/Groups frame. Once the frame is expanded, a Restore button appears for collapsing the frame.

The filter options menu provides the following options.

Display Menu

Provides the following options for displaying and selecting the filter values.

Dropdown List

Generates a drop-down list for filter values. When expanded, displays a checklist of values.

Check List

Lists the field values with a check box next to each, and a scroll bar, if needed to scroll through the values. This is the default display option for filters on alphanumeric and integer fields.

Slider

Available for filters on non-integer numeric fields only. Displays a bar with two sliders for selecting From and To values. When you move the sliders, a Reset button displays that resets the range to the entire range of values. This is the default display option for filters on non-integer numeric fields.

Text Box

Adds a text box in which to enter a character string or pattern that uses the wildcard characters supported with the LIKE operator. A percent symbol (%) in the pattern means that any number of characters in that position matches the pattern (added automatically at the end of the character string you enter). An underscore character (_) means that a single character in that position matches the pattern. A plus sign on the menu bar enables you to add additional values, and a check mark enables you to apply the filter.

Double List

Shows two lists, one with available values and one with selected values. To add values to the Selected Values list, you can right-click a value or multi-selected values (if Multiple Values is enabled) and choose Select. You can also double-click a value or drag values. To remove values from the Selected Values list, you can right-click a value or multi-selected values and choose Unselect. You can also double-click a value or drag values.

Once you have added a value to the Selected Values list, an *Unselect all* button appears to make it easy to remove all selections.

Multiple Values

Enables multiple selections. This is the default selection option.

Single Value

Disables multiple selections.

Plain

Displays a one-column list of values. This is the default display.

Two columns (drop a field)

Enables you to drop a field that contains a caption (description) for each filter value in a second column.

Relation Menu

Provides the following operators for the filter.

EQ

Selects fields with values equal to a filter value. This is the default operator for alphanumeric and integer fields.

NE

Selects fields with values not equal to any filter value.

LT

Selects fields with values less than the filter value.

GT

Selects fields with values greater than the filter value.

LE

Selects fields with values less than or equal to the filter value.

GE

Selects fields with values greater than or equal to the filter value.

LIKE

Selects fields whose values match a pattern. This operator is available for alphanumeric fields where the display is a text box.

RANGE

Selects fields whose values fall within a range. This operator is available for numeric fields. This is the default operator for non-integer numeric fields.

Variables Menu

Supported for filters on alphanumeric fields. It lists all defined variables for selection and includes the default option to prompt for actual data.

Values Menu

Supported for filters on alphanumeric fields. Provides the following options.

Show Null

Shows null values on the list of selection values.

Hide Nulls

Does not show null values on the list of selection values. This is the default option.

Convert to &&-variable

Supported for a check box filter that uses the EQ or NE relation and has a limited number of values. Converts the filter to a global variable expression that improves performance.

Lookup other file

Enables you to create a DB_INFILE expression that checks each value against values that exist in another file.

Advanced Menu

Provides the following options.

No delay, immediate update

Updates the filter as soon as you select a value. This is the default update option.

Use Apply button

Does not update the filter until you click the *Apply Filter* button, which displays on the menu bar when you select this option.

Show selected values on top

Moves selected values to the top of the list. This is the default option.

Keep selected values in place

Does not move selected values to the top of the list.

Delete Filter

Deletes the filter from the synonym and the Filters/Groups frame.

Properties

Opens a dialog box showing filter properties in which you can change the title for the filter that appears on output.

Group Menu

The down arrow (menu) for a group provides the following options.

Display Menu

Provides the following options for displaying groups.

Double List

Shows two lists, one with input values and one with selected values. To add values to the Selected Values list, you can right-click a value or multi-selected values and choose Select. You can also double-click a value or drag values. To remove values from the Selected Values list, you can right-click a value or multi-selected values and choose Unselect. You can also double-click a value or drag values. If you right-click in the Input Values list and select *Make a new group*, the new group will be added in the Selected Values list. The added groups are named based on the values selected and contain the values selected.

[Single List](#)

Shows a list of values, initially in one default group. Right-clicking the menu options let you create a new group, copy and paste values, ungroup values, edit the group name, and recreate groups based on like-sounding values.

[Add All](#)

Adds all shown ungrouped values to a new group.

[Create a New Group](#)

Creates a new group with no values in it.

[Recreate groups by Metaphone](#)

Uses the Metaphone encoding algorithm to group alphanumeric values by pronunciation.

[Recreate groups by Soundex](#)

Uses the Soundex encoding algorithm to group alphanumeric values by pronunciation.

[Ungroup All](#)

Removes all groups except the default group.

[Edit properties](#)

Opens the Edit properties dialog box. You can change the overall group name, title, and default value.

[Hide Group Card](#)

Removes the group from the frame, but leaves it in the synonym. To reopen the group in the Filters/Groups frame, right-click the group in the synonym and select *Edit Grouping*.

Reference: Decode Pane

A decode allows you to specify alternate values for each value of a field using the DECODE function.

To create a decode, right-click a field, click *New Expression*, then *Decode*.

Each decode added opens a card where you can specify original and associated values. For example, the Decode shown in the following image associates text values with numeric gender values.

Default Value: ✎

Code ↓	Result ↓
0	<input type="text" value="undefined"/>
1	<input type="text" value="male"/>
2	<input type="text" value="female"/>

You can specify a default value for those entries for which you do not enter a value.

The Edit menu at the top allows you to change the field name and title.

Click outside the Decode card to add the new field.

Reference: Output Pane

The Output pane shows the results of changes to a synonym or view, sample data output, and output of sample data, impact analysis, and data profiling reports and charts.

Each new output can be added on a new tab in the Output pane or replace the existing tab, depending on your preferences. The following image shows the output pane with multiple tabs. If there are too many tabs to display in the window, left/right arrows (< >) display for scrolling.

ID	Sal	Cost of Goods, Local Currency	Cost of Goods, Gro	Cost of Gooc	Discount, Local Curren	Disc
1		\$234.00	\$234.00	\$234.00	\$0.00	\$0.00
2		\$46.00	\$46.00	\$46.00	\$20.00	\$20.00
3		\$380.00	\$380.00	\$380.00	\$49.00	\$49.00
4		\$374.00	\$374.00	\$374.00	\$0.00	\$0.00
5		\$310.00	\$310.00	\$310.00	\$0.00	\$0.00
6		\$83.00	\$83.00	\$83.00	\$0.00	\$0.00
7		\$312.00	\$312.00	\$312.00	\$71.70	\$71.70
8		\$548.00	\$548.00	\$548.00	\$0.00	\$0.00
9		\$400.00	\$400.00	\$400.00	\$0.00	\$0.00
10		\$131.00	\$131.00	\$131.00	\$0.00	\$0.00
11		\$243.00	\$243.00	\$243.00	\$131.39	\$131.39
12		\$370.00	\$370.00	\$370.00	\$0.00	\$0.00
13		\$135.00	\$135.00	\$135.00	\$57.50	\$57.50
14		\$374.00	\$374.00	\$374.00	\$150.00	\$150.00
15		\$16.00	\$16.00	\$16.00	\$0.00	\$0.00
16		\$240.00	\$240.00	\$240.00	\$0.00	\$0.00
17		\$390.00	\$390.00	\$390.00	\$90.00	\$90.00
18		\$165.00	\$165.00	\$165.00	\$0.00	\$0.00
19		\$410.00	\$410.00	\$410.00	\$0.00	\$0.00
20		\$83.00	\$83.00	\$83.00	\$0.00	\$0.00
21		\$108.00	\$108.00	\$108.00	\$0.00	\$0.00
22		\$80.00	\$80.00	\$80.00	\$0.00	\$0.00

The menu bar at the top of the pane provides the following options.

[Double the row limit](#)

Doubles the number of rows used for sample and changed data display.

The View menu (down arrow) provides the following options.

[Choose Columns ...](#)

Opens a dialog box for selecting the metadata attributes and number of field values to display.

[Reset to Defaults](#)


Resets the pane to the default view.

On each tab, each column title has a menu (opened using the down arrow) that has the same options as the right-click menu for the field in the Table/Column and Business View panes. For information, see [Right-Click Options for Fields in the Table/Column, Business View, and Output Panes](#) on page 321.


In addition, if you click a column heading, the sort order for that tab toggles between sort ascending, sort descending, and original order of the column in which you clicked.

Reference: Join Frame

The Join frame shows a pictorial view of the joins between tables in a cluster, as shown in the following image.



If you use the Flow view, the synonyms in the join are represented by database-specific icons, as shown in the following image.



You can edit the joins by adding and deleting tables and by changing the join properties. For more information, see [Creating a Join](#) on page 308.

Enhancing a Synonym

You can enhance a synonym by:

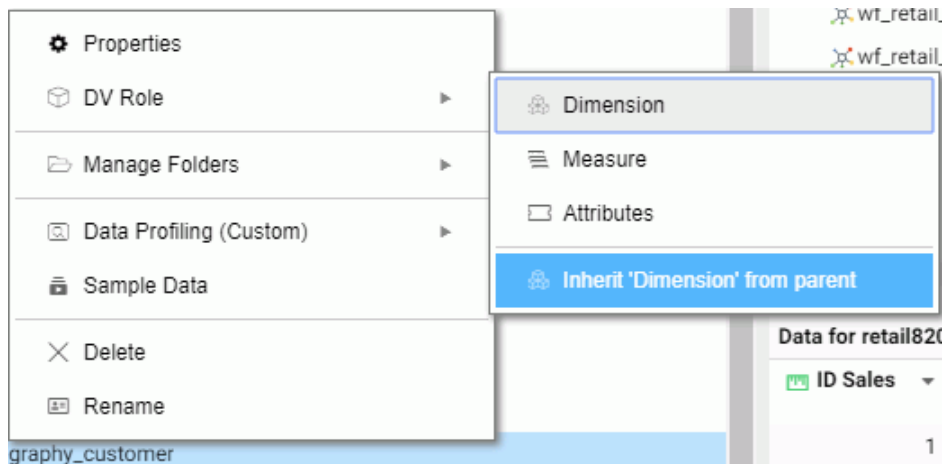
- Assigning DV roles. For information, see [Assigning Dimension View \(DV\) Roles](#) on page 297.
- Adding filters. For information, see [Right-Click Options for Fields in the Table/Column, Business View, and Output Panes](#) on page 321.
- Creating a virtual field or calculated summary value. For information, see [Creating a Virtual Field \(DEFINE\) or Calculated Value \(COMPUTE\)](#) on page 299.
- Adding a global variable. For information, see [Adding a Global Variable to a Synonym](#) on page 306.
- Creating a join. For information, see [Creating a Join](#) on page 308.
- Adding security restrictions. For information, see [Adding Data Source Security \(DBA\)](#) on page 314.
- Pivoting columns. For information, see [Pivoting Columns](#) on page 319.

Note: When you enhance a synonym, in most cases you can use the Business View frame to create the enhancement. However, since this frame gives you a logical view of the synonym, it does not give you full access to the features that are added. For example, you can hide a virtual field or calculated value so that it does not appear in the Business View, but that will not delete it from the synonym. To have a full view of and access to synonym components, you must open the Table/Column frame, if it is not already open. To open the Table/Column frame, click *View*, then *Show "Table/Column"*. This frame represents the physical view of the synonym and provides full access to all of its components.

Assigning Dimension View (DV) Roles

The Business View frame defines folders that function as segments to provide a view of the synonym and to define the accessible fields and their relationships. Folder relationships are the same as segment relationships, with parent folders, child folders, and sibling folders.

You can assign a DV role to a folder or field by right-clicking the folder or field and selecting a DV role, as shown in the following image.



You can explicitly assign a DV role to a folder or field, or have it automatically inherit its role from its parent. If you explicitly assign a DV role, that role moves with the object if you drag it to another location within the structure. If you do not explicitly assign a DV role, the role changes as you move the object under a new parent, except if you drop it onto a field with the Drill Level role. If dropped onto a Drill Level field, the moved field inherits the Drill Level role.

The following DV roles can be assigned.

- ❑ **Dimension.** A dimension field, when double-clicked or dragged onto the report or chart canvas in the WebFOCUS tools, will automatically be added to the request as a vertical (BY) sort field.

A folder can be assigned the role Dimension.

A field can be assigned the role Dimension (Standalone) or Dimension (Drill Level). When it is assigned the role Dimension (Drill Level), it will become part of a hierarchy where the levels depend on the order of the fields in the folder. Then, when AUTODRILL is turned on, automatic drill-downs will be created on the report or chart output.

For a folder assigned the DV role Dimension or a field assigned the DV role Dimension (Standalone), the following attribute is added to the folder or field declaration in the synonym.

`DV_ROLE=DIMENSION`

For a field assigned the DV role Dimension (Drill Level), the following attribute is added to the field declaration in the synonym.

`DV_ROLE=LEVEL`

A folder can contain only one drill level hierarchy. However, you can use the same fields in multiple hierarchies by placing each hierarchy in a separate folder. A folder with a drill level hierarchy is not limited to just the hierarchy. It can contain other fields with different DV_ROLES, although this is not recommended.

- ❑ **Measure.** A measure field, when double-clicked or dragged onto the report or chart canvas in the WebFOCUS tools, will automatically be added to the request as an aggregated value (SUM), if it is numeric. If it is alphanumeric, it will be added as a vertical (BY) sort field. A folder or field can be assigned the role Measure.

For a folder or field assigned the DV role Measure, the following attribute is added to the folder or field declaration in the synonym.

`DV_ROLE=MEASURE`

- ❑ **Attribute.** An attribute field, when double-clicked or dragged onto the report or chart canvas in the WebFOCUS tools, will automatically be added to the request as a vertical sort field (BY).

For a folder or field assigned the DV role Attribute, the following attribute is added to the folder or field declaration in the synonym.

`DV_ROLE=ATTRIBUTE`

- ❑ **Folder.** A folder is a virtual segment in a BV. It can be assigned the roles Dimension, Measure, Attribute, or Folder.

Note: When a folder is inserted as a child of a field, the attribute PARENT_FIELD describes this relationship. By default, such a folder and its fields will be assumed to have the Attribute role.

- ❑ **None.** If no role is assigned, the field or folder will inherit its role from its parent. If a role has been assigned, you can remove it by selecting the option to inherit its role from its parent.

Example: Sample Dimension Folder Declaration

The following declarations show sample folder and field definitions. Note that the declaration for each field in a folder specifies the real segment it actually belongs to.

The DV_ROLE for the PRODUCT folder is DIMENSION. The DV_ROLE for the PRODUCT_CATEGORY field is LEVEL.

```
FOLDER=PRODUCT, PARENT=DIMENSIONS,
DV_ROLE=DIMENSION,
DESCRIPTION='Product', $
FIELDNAME=PRODUCT_CATEGORY, BELONGS_TO_SEGMENT=WF_RETAIL_PRODUCT,
DV_ROLE=LEVEL, $
FIELDNAME=PRODUCT_SUBCATEG, BELONGS_TO_SEGMENT=WF_RETAIL_PRODUCT,
DV_ROLE=LEVEL, $
FIELDNAME=MODEL, BELONGS_TO_SEGMENT=WF_RETAIL_PRODUCT,
DV_ROLE=LEVEL, $
```

Creating a Virtual Field (DEFINE) or Calculated Value (COMPUTE)

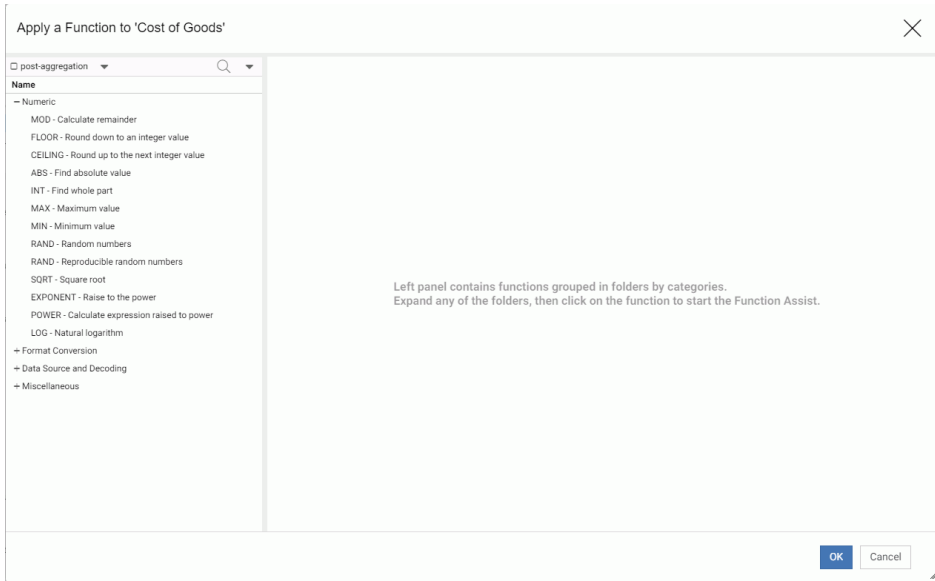
A virtual field can be used in a request as though it is a real data source field. The calculation that determines the value of a virtual field is performed on each retrieved record that passes any screening conditions on real fields. The result of the expression is treated as though it were a real field stored in the data source. A virtual field defined in a synonym is available whenever the data source is used for reporting.

A calculated value is a temporary field that is evaluated after all the data that meets the selection criteria is retrieved, sorted, and summed.

Right-click a field, point to *New Expression*, and click:

- ❑ *Apply Function*, to apply a function to the field.

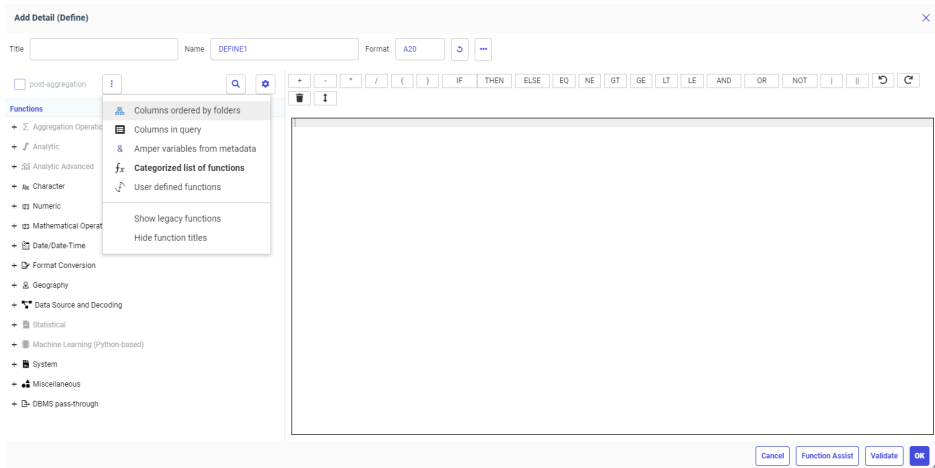
The Apply a Function dialog box opens, as shown in the following image.



When you have entered the required information, click **OK**. For more information, see [Other Synonym Enhancements](#) on page 321.

Advanced Expression to create a complex expression.

The Advanced Expression calculator opens, as shown in the following image.



For functions, by default, the post-aggregation check box is not selected. This creates a DEFINE. If you click the check box, you will create a COMPUTE.

Create the Expression and click *OK*.

For more information about creating an expression see [Using the Expression Calculator](#) on page 301.


- Group and Replace** or **Phonetic Grouping**. Opens the Filters/Groups frame and creates a group field using values of the selected field. For more information, see [Filters/Groups Frame](#) on page 287 and [Other Synonym Enhancements](#) on page 321.
- Decode**. Opens a wizard for associating a new value with each existing value of the field.
- Create Bins**. Opens a wizard for dividing the data values in the field into a specified number of bins based on equal intervals, equal intervals using a formula, or quantiles. For more information, see [Other Synonym Enhancements](#) on page 321.

Using the Expression Calculator

When you generate a virtual field or calculated value, the Expression Calculator opens with buttons and fields that contain all of the elements needed for generating any type of WebFOCUS expression.

The expression calculator is used for defining virtual fields (DEFINE) and creating calculated summary values (COMPUTE).

To open the Expression Calculator, right-click a field, point to *New Expression*, and click *Advanced Expression* on the context menu. The Expression Calculator is shown in the following image.



Top Pane Options

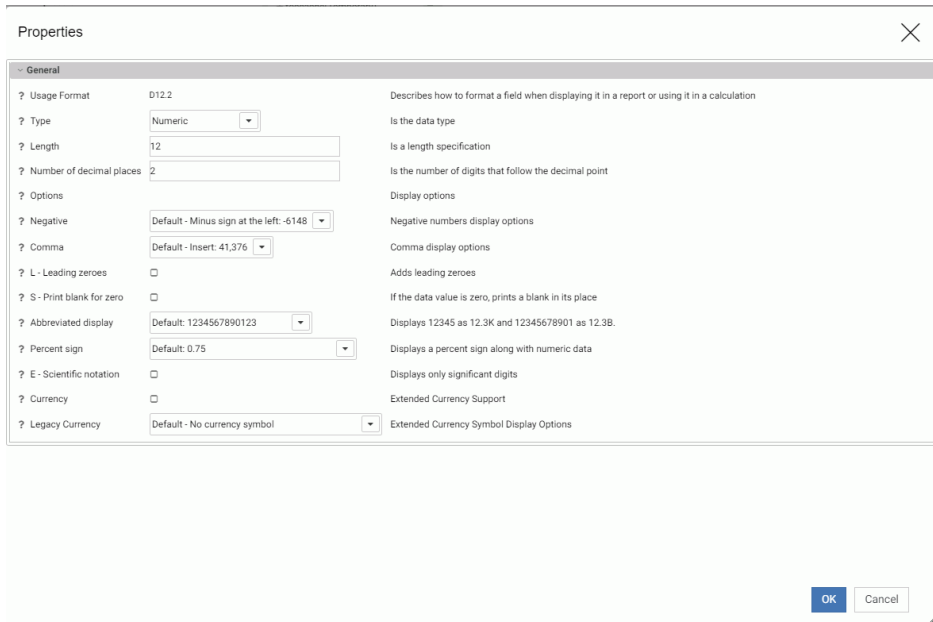
Enter a Title for the virtual field or calculated value in the Title text box, or leave it blank. If you enter a title, the default Name will be changed to be the title, with special characters converted to underscores when you click outside the Title box. Alternatively, you can edit the default field name.

The title will be used as the column heading on reports, if you enter one. If you do not, the field name will be used.

Enter a format in the Format text box, accept the default (A20), click *Refresh* to recalculate the format, or click the ellipsis (...) to open the Properties dialog box.

You must make sure the format supports the expression you will be creating.

The Properties dialog box is shown in the following image, with Numeric selected from the Type drop-down list.



The following is the list of data types.

Data Type	Default Usage Format
Character	A20V

Data Type	Default Usage Format
Integer	I4
Numeric	D12.2
Date	YYMD (order of components depends on the Locale settings)
Date and Time	HYYMDS (order of components depends on the Locale settings)
Character (Fixed)	A20
Text	TX256
Decimal	P8.2
Time	HHIS
Hexadecimal	U20
Binary Large Object	BLOB

The display options applicable to the selected data type display. As you select options, the Usage format at the top of the window changes to show the selected options. For more information, see the *Describing Data With TIBCO WebFOCUS® Language* manual.

View Options

The View options list menu provides the following options.

- List of functions.** The function categories that display initially are the recommended categories.
- post-aggregation.** Clicking the post-aggregation check box adds function categories that can be applied to fields after they have been aggregated. This also changes the calculation from a DEFINE to a COMPUTE.
- Columns ordered by folders.** Displays fields, filters, and variables in the Business View instead of a list of functions.
- Amper variables from metadata.** Displays the list of global variables referenced in the synonym.

- Categorized list of functions.** This is the initial view. You can return to this view with this button.
- Advanced Options.** Provides additional view options:
 - When the view is **Categorized list of functions.**, enables you to select *User Functions*, which displays the list of applications in your APPPATH so you can navigate to your own functions, *Show legacy functions*, to add legacy functions to the display, or *Hide function titles* to display function names without descriptive text. If you select *Show legacy functions*, that option changes to *Hide legacy functions*. If you select *Hide function titles*, that option changes to *Show function titles*.

Note: To be recognized as a procedure that contains user functions, any .fex file that contains user functions (DEFINE FUNCTIONS) must have the following syntax at the beginning.

```
-*DM_JOB_TYPE=512
```
 - When the view is **Columns ordered by folders**, enables you to change to *Tables/Columns*, which displays the physical view of the synonym instead of the Business View, or *User Functions*, which displays the list of applications in your APPPATH so you can navigate to your own functions.

When you have chosen a view option, the option you have chosen is used as the title above the pane of items listed for that option.

Expression Pane Options

In the Expression pane, you can enter an expression in the text box or double-click fields from the metadata pane and use the calculator buttons to create the expression. For more information about expressions and operators, see the *Creating Reports With TIBCO WebFOCUS® Language* manual. The buttons are:

- Relational Operators.**
 - LT** (Less than). Is true if the expression on the left evaluates to be less than the expression on the right.
 - EQ** (Equal to). Is true if the expression on the left evaluates to be equal to the expression on the right.
 - GT** (Greater than). Is true if the expression on the left evaluates to be greater than the expression on the right.
 - LE** (Less than or equal to). Is true if the expression on the left evaluates to be less than or equal to the expression on the right.

- NE** (Not equal to). Is true if the expression on the left evaluates to be not equal to the expression on the right.
- GE** (Greater than or equal to). Is true if the expression on the left evaluates to be greater than or equal to the expression on the right.
- Conditional Operators.**
 - IF** (specifies the condition to be tested).
 - THEN** (specifies the result if the condition is true).
 - ELSE** (specifies the result if the condition is false).
- Logical Operators.**
 - AND** (is true if both expressions are true).
 - OR** (is true if at least one of the expressions is true).
 - NOT** (is true if the expression is false).
- Alphanumeric Operators.**
 - |** (weak concatenation).
 - ||** (strong concatenation, remove trailing blanks).
- Arithmetic Operators.**
 - **** (raise to a power).
 - +** (plus). Add the numbers on either side of the + sign.
 - (minus). Subtract the number on the right of the minus sign from the number on the left.
 - *** (multiplication). Multiply the number on the left by the number on the right.
 - /** (division). Divide the number on the left by the number on the right.
- Undo.**

Undoes successive actions.
- Redo.**

Redoes successive actions.

Clear all (trash can).

Clears the expression.

Show buttons at bottom (double up and down arrow).

Moves the buttons to the bottom of the pane, and changes to *Show buttons on top*.

Other Buttons.

Validate checks if the expression entered is valid.

Function Assist helps you apply a function to an expression. First, you must change the view in the metadata pane to show the list of functions and select a function. Clicking *Function Assist* will open a dialog box that describes what the function does and allows you to select fields for the parameters from the metadata. Function Assist opens automatically if you double-click a function name. You can click the *Example* button to see an example of the function syntax. Click *OK* to add the function to the expression or *Cancel* to cancel.

When you have finished creating the expression, click *OK* to add the virtual field or calculated value to the metadata.

Adding a Global Variable to a Synonym

You can define global variables in a Master File, filter, or expression and use them to parameterize certain attributes in the Master File and its corresponding Access File. You can also add certain system variables to a synonym.

To insert a global variable:

1. In the Table/Column frame, right-click the synonym name and click *New Variable* on the context menu, or in the Business View pane, right-click the synonym name, click *Insert*, then *Variable*.

The Properties window opens, as shown in the following image.

The screenshot shows the 'Properties' dialog box with the 'General' tab selected. The 'Variable Name' field contains '&&VARIABLE1'. The 'Prompt Values' dropdown menu is set to 'Input box'. The 'Persistent' checkbox is checked. The 'Usage Format' dropdown menu is empty. The 'Advanced' tab is also visible below the 'General' tab.

2. Enter a name for the variable or accept the default.
3. Enter a prompt caption, if needed, for describing what type of value to enter when prompted for the value.
4. Enter a default value, if there is one.
5. Select the type of value container to show, when a request is run.

You can select from the following General options:

- Input box.** The user must enter a value.
- Static list of values.** The user must select from a predefined list of values. If you select this option, additional input fields appear in which you list the static values and select whether the user can select a single value or multiple values.

Enter a list of values enclosed in single quotation marks and separated by the logical operator OR.

- Static list of values/captions.** The user must select from a predefined list of values. If you select this option, additional input fields appear in which you list the static values and the captions that should display instead of the actual values, and select whether the user can select a single value or multiple values.

Enter a list of pairs, each pair containing the actual value and the caption to appear separated by commas. The list must be enclosed in parentheses. Each entry must be enclosed in single quotation marks and each pair must be separated by commas.

- Synonym based dynamic list of values.** The user must select from a list of values that is generated dynamically by looking up field values in a lookup synonym. If you select this option, additional fields appear for entering or selecting the lookup synonym, the lookup field, the display field for the caption to appear, and whether the user can select only a single value or multiple values.
- Procedure based dynamic list of values.** The user must select from a list of values that is generated dynamically by running a procedure that retrieves the lookup field values and display field values. If you select this option, additional fields appear for entering or selecting the lookup procedure, any parameters required for running the procedure, the lookup field, the display field for the caption to appear, and whether the user can select only a single value or multiple values.

You can also enter or select the following Advanced properties for the variable.

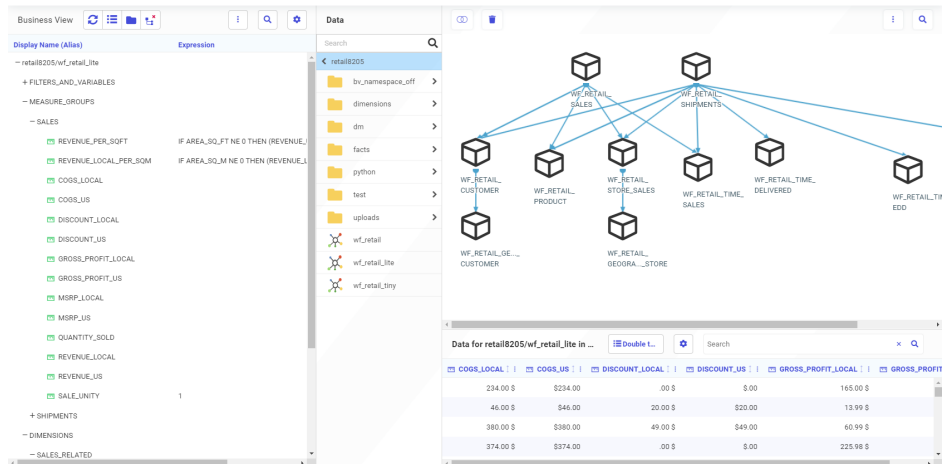
- Whether to add single quotation marks around strings.
 - Whether to allow the default value assigned to the variable to persist across calls to WebFOCUS or to be available to change. When allowed to change, the variable can be used as a selection filter in WebFOCUS Designer.
 - The data type.
 - The length of the variable, if you select a data type.
6. Click *Apply* to save these properties or *Cancel* to cancel.

Creating a Join

A join describes a relationship between multiple synonyms based on a condition. The condition may be equality between fields in the synonyms or the result of an expression that relates fields in the synonyms. For more information about joins in a synonym, see the *Describing Data With TIBCO WebFOCUS® Language* manual.

To create a new cluster synonym, from the Applications page either right-click an application, point to *New*, and click *Cluster Business View*, or click *New* on the ribbon, then *Cluster Business View*. To edit an existing join, open an existing synonym in Data Assist.

The Synonym Editor opens, as shown in the following image for an existing synonym.



If you are creating a new cluster synonym, you can double-click or drag synonyms from the Data frame tree into the Join Editor.

You can also make the following edits from the Join frame menu bar:

Options

Sets the following join defaults.

- Insert child segment with snowflakes.** The default is No, to insert only the segment selected.
- Display full component names.** The default is Yes.

Click *OK* to save your defaults.

View

The View options for Data Steward (Model) Mode are:

- Collapse All.** Collapses all of the synonyms so that the fields do not show.
- Expand Referenced.** Expands the referenced synonym icons to show four fields.
- Expand All.** Expands all of the synonym icons to show four fields.
- Set to Standard Size.** Shows four fields at a time.
- Expand to Full Size.** Expands the synonym boxes to full size.
- Expand to Custom Size.** Options are Height Standard, Height Full, Width Standard, and Width Full.

- Auto Arrange.** Automatically arranges the synonyms.

The View options for Analyst (Flow) Mode are:

- Layout.** Click Vertically spread out (the default), Vertically Compressed, Vertically Compressed with long H-Space, and Default.
- Auto Arrange.** Automatically arranges the synonyms and joins.

Mode

Displays the joins in either Data Steward Mode (the default) or Analyst Mode. In Analyst mode, the synonyms are represented by an icon, and each join is represented by a container that displays the type of join. In Data Steward mode, each synonym is represented by a cube that you can double-click or set the View option to Expand All in order to display field names, and each join is represented by a vector between the joined synonyms.

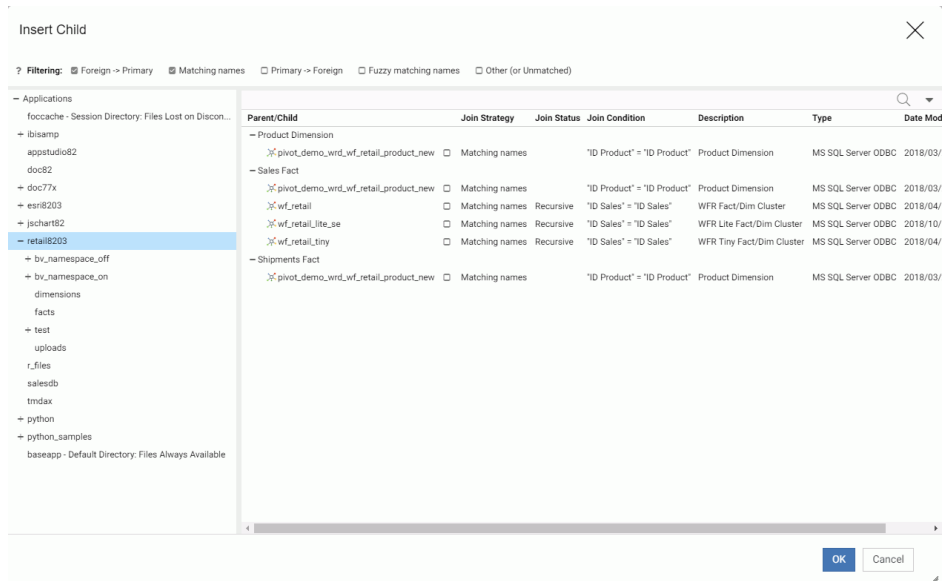
Delete All Joins

Deletes all joins.

Insert

You can insert a child or a root in Data Steward mode. You can add a join in Analyst mode (provides the same options as Insert Child in Data Steward mode).

- Insert Child.** The Insert Child dialog box opens, as shown in the following image.



Select a synonym to insert as a child by selecting the check box to the right of the synonym name.

Use the Filtering check boxes to help with selection. You can filter the synonym selections using one or more of the following filtering options:


- Foreign to primary key matches.
- Matching names.
- Primary to foreign key matches.
- Fuzzy matching names (not exact matches).
- Other (or unmatched).

The View menu enables you to view the synonyms as icons or view the details. You can also use it to choose columns and reset the view to the default view.

The Find icon (magnifying glass) opens a search box.

Click *OK* when you are finished.

- Insert Root** The Insert Root dialog box opens, as shown in the following image.



Select a synonym to insert as a root by selecting the check box to the right of the synonym name in the More pane.

Use the Filtering menu to select one of the following filtering options:

- All tables.
- Tables with foreign keys.

Click *OK* when you are finished.

You can edit the join by right-clicking a synonym icon or connector line.

Editing a Join by Right-Clicking a Synonym Icon

From a synonym icon, select from the following options.

- Manage Parent Links.** Only available in Data Steward mode for a child synonym. Will say *Edit Parent Link* if there is only one parent. Opens a dialog box for selecting a parent synonym. Right-clicking a parent and clicking *Manage* opens the Edit Parent Link dialog box for selecting new join fields and changing the join type and condition.

For more information about editing a parent link, see [Editing a Parent Link](#) on page 313.

- Insert Child.** Opens the Insert Child dialog box. Not available in Analyst mode.
- Add/Remove Parent Links.** Opens the Add/Remove Parent Links dialog box. You can check and uncheck synonyms to add or remove parent links. Not available in Analyst mode.
- Sample Data.** Generates sample data.
- Sample Data with parent key.** Generates sample data that includes the key from the parent synonym.
- Data Profiling (Custom).** Choose from Statistics, Key Analysis (Data Steward mode only), and Count.
- Delete.** Deletes the join.

Editing a Join by Right-Clicking a Connector Line

From a connector line, click one of the following options:

- Edit Parent Link.** Opens the Edit Parent Link dialog box for editing the join. In Analyst mode, is available as Edit Link when right-clicking a join icon.

For more information about editing a parent link, see [Editing a Parent Link](#) on page 313.

- Join Multiple.** Makes the join a one-to-many join.
- Join Unique.** Makes the join a one-to-one join.
- Inner Join.** Makes the join an inner join.
- Left Outer Join.** Makes the join a left outer join.
- Right Outer Join.** Makes the join a right outer join.
- Full Outer Join.** Makes the join a full outer join. Full outer joins are supported even if the relational data source does not support them.

- Delete Link.** Deletes the join.
- Data Profiling (custom).** Produces data profiling reports.
- Sample Data with parent key.** Produces sample data.

Editing a Parent Link

The Edit Parent Link dialog box in the Reporting Server browser interface is shown in the following image.

The screenshot shows the 'Edit Join from Store Sales Dimension to Store Sales Geography Dimension' dialog box. The left pane shows the 'Configure Join' section with radio buttons for 'Inner', 'Left Outer', 'Right Outer', and 'Full Outer'. The 'Left Outer' option is selected. Below this, there are fields for 'ID Geography' and 'Store ID Geography' with a '+' button and a 'Suggestions' button. The right pane has tabs for 'Sample Data' and 'Join profiling'. The 'Sample Data' tab is active, showing a table with columns for 'Sales Fact ID Store', 'Store Sales Dimension ID Store', 'Store Sales Dimension ID Geography', 'Store Sales Geography Dimension Store ID Geography', 'Sales Fact ID Sales', 'Sales Fact ID Currency', 'Sales Fact ID Customer', 'Sales Fact ID Discount', and 'Sales Fact ID'. The table contains 17 rows of data.

Sales Fact ID Store	Store Sales Dimension ID Store	Store Sales Dimension ID Geography	Store Sales Geography Dimension Store ID Geography	Sales Fact ID Sales	Sales Fact ID Currency	Sales Fact ID Customer	Sales Fact ID Discount	Sales Fact ID
0	0	84506	84506	1	1	173889	100	
0	0	84506	84506	2	1	245541	104	
0	0	84506	84506	3	1	230444	101	
0	0	84506	84506	4	1	213958	100	
0	0	84506	84506	5	1	23113	100	
0	0	84506	84506	6	1	232043	100	
0	0	84506	84506	7	1	11162	102	
0	0	84506	84506	8	1	200808	100	
0	0	84506	84506	9	1	61715	100	
0	0	84506	84506	10	1	55649	100	
0	0	84506	84506	11	1	159686	105	
0	0	84506	84506	12	1	28928	100	
0	0	84506	84506	13	1	56206	104	
0	0	84506	84506	14	1	40414	104	
0	0	84506	84506	15	1	189952	100	
0	0	84506	84506	16	1	72393	100	
0	0	84506	84506	17	1	238885	103	


The left frame displays the join conditions and editing options. The right pane has tabs for sample data reports and join profiling reports and charts.

Configure Join Options.

Click *Join Type* to select one of the following join types:

- Inner Join.** Selects only those records that have values that satisfy the join condition in both synonyms.
- Left Outer Join.** Selects all records from the host synonym. Where there is no record that satisfies the join condition in the target file, default values are substituted.
- Right Outer Join.** Selects all records from the target synonym. Where there is no record that satisfies the join condition in the host file, default values are substituted.
- Full Outer Join.** Selects all records from the host and target synonyms. Where there is no record that satisfies the join condition in the host or target file, default values are substituted.

Click **Multiple.** to make the join a one-to-many join.

In the Data Migrator desktop interface, you can also click the Add New Row button () to add a new row in the Join Condition section of the editor and add a new expression to the join condition.

The View menu provides the following options.

- Show title (the default).
- Show name.
- Show source name.
- Show sample values.

Click the host or target field to select a different field. Click the pencil icon to open a double list for selecting both.

Click the down arrow next to the relation to change the relational operator.

Click the plus sign (+) under the host field to add a new join condition.

Click *Expression* to add an expression using the Expression Calculator.

Click *Clear All* to clear all conditions.

Click *Suggestions* to open a pane with potential join conditions. Clicking the plus sign (+) next to a condition adds it to the join. Clicking *Suggestions* again closes the pane.

Click *Clear All (X)* to remove all of the join conditions.

On the Sample Data tab:

- Click *Find* to open a search box.
- Click *View* to choose columns or set the view back to the default.

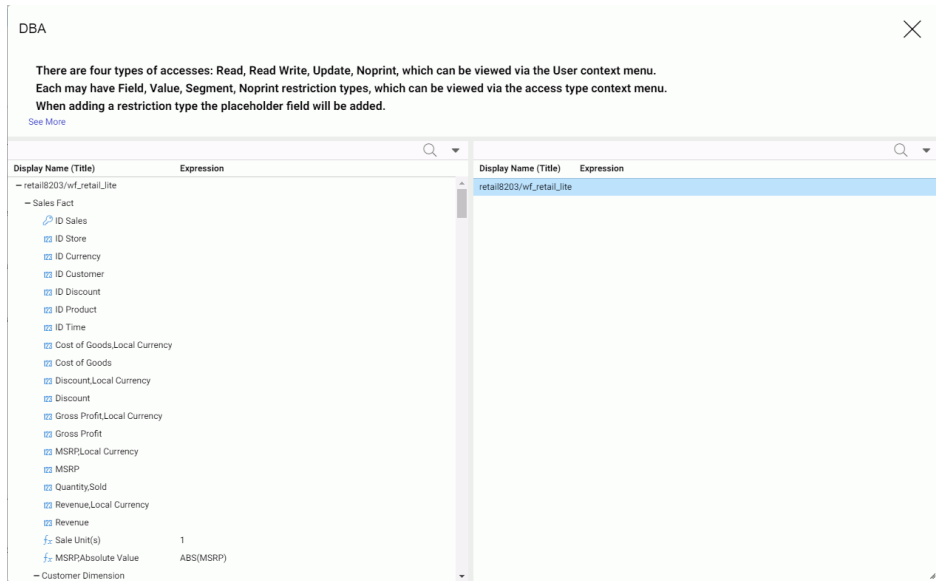
Adding Data Source Security (DBA)

The DBA facility allows you to define restrictions and access rights for users.

Adding DBA Security

1. Click *DBA* on the *Advanced Options* menu of the ribbon.

The DBA window opens, as shown in the following image.



2. In the right pane, right-click the synonym name, point to *Insert*, and click *DBA*.

- The top node, named DBA1 by default, is the password for the Database Administrator (DBA). The DBA has complete access, so no access rights need to be defined.
- The next node, named USER0001 by default, represents a user password. By default, the user is given read access to the data source.

You can rename the passwords and add and change access rights using right-click options. You can undo actions using the *Undo* link on the menu bar. Once you have created DBA attributes, you can delete the DBA section of the Master File by Clicking *Remove DBA* on the menu bar.

3. To change a password, right-click a user and click *Rename*.

A *Rename* dialog box opens.

Enter the new password and click *OK*.

4. To add a user, right-click the DBA password, point to *Insert*, and select *User*.

A new user node is added to the DBA Statements pane.

5. To add an access right to a user password, right-click a user password, point to *Insert*, and click one of the following:

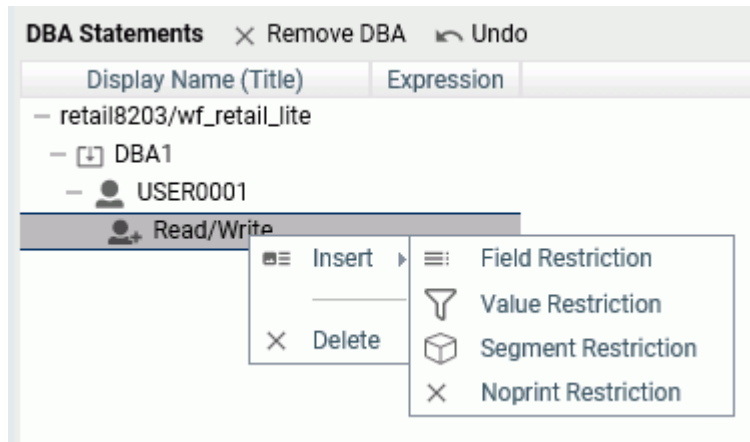
- Read Access.** This allows the user to only read the data source.
- Write Access.** This allows the user to only write new segment instances to the data source.
- Read/Write Access.** This allows the user to both read from and write to the data source.
- Update Access.** This allows the user to only update records in the data source.

The access right is added under the user password. You can then restrict those rights to specific segments, fields, or field values.

6. You can further restrict the access rights of a user based on segment names, field names, and field values.

- a. **Field Restriction.** To prevent access rights to a field, do the following:

- a. Right-click an access right under the user you want to restrict, point to *Insert*, and select *Field Restriction* as shown in the following image.



A field restriction placeholder is added under the access right.

- b. Drag a field to the placeholder.

The placeholder is replaced by the field.

- c. To restrict access to additional fields, repeat these steps for each additional field.

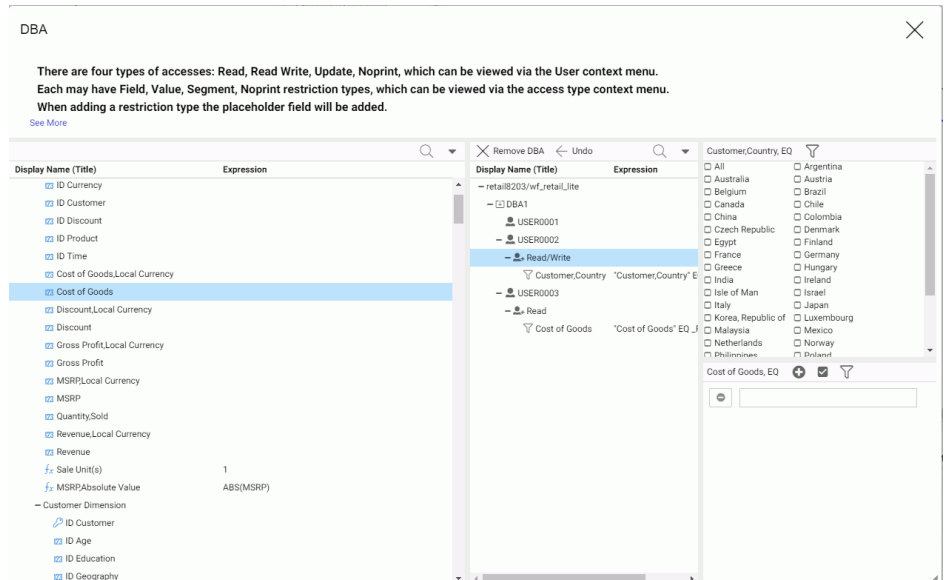
- b. **Value Restriction.** To prevent access rights to specific field values, create a filter.

- a. Right-click an access right under the user you want to restrict, point to *Insert*, and select *Value Restriction*.

A value restriction placeholder is added under the access right.

b. Drag a field to the placeholder.

The placeholder is replaced by the field, with a default expression. A Restrictions dialog box opens for you to create the expression that will identify the restricted values, as shown in the following image.



For a non-integer numeric field, a text box opens in which you can enter values. Using the menu on the top right of the Restrictions pane, you can change the display to a slider or a checklist. For alphanumeric or integer fields, the default display is the checklist. You can multi-select values for the EQ and NE relations. If your display option is a text box, you can click the Add Value (+) button on the menu bar to add another text box. To delete a value, click the Delete Value (-) button on the left of the text box that contains the value you want to delete.

By default, the EQ (equal to) relation is used, so the restriction is in effect for field values equal to the values you enter or select. To change the relation, point to Relation on the menu and select one of the following relational operators.

- EQ.** Equal to. The restriction is in effect for field values equal to any of the values specified in the restriction expression.
- NE.** Not equal to. The restriction is in effect for field values not equal to any of the values specified in the restriction expression.

- LT.** Less than. The restriction is in effect for field values less than the value specified in the restriction expression.
- GT.** Less than. The restriction is in effect for field values greater than the value specified in the restriction expression.
- LE.** Less than or equal to. The restriction is in effect for field values less than or equal to the value specified in the restriction expression.
- GE.** Less than. The restriction is in effect for field values greater than or equal to the value specified in the restriction expression.
- Range.** From-to. The restriction is in effect for field values greater than or equal to the first value in the range and less than or equal to the second value in the range.

Once you have created an expression, additional options become available on the menu bar. You can update the expression using the Apply button.

- c. **Segment Restriction.** To prevent access to one or more segments, do the following:
 - a. Right-click an access right under the user you want to restrict, point to *Insert*, and select *Segment Restriction*.

A segment restriction placeholder is added under the access right.
 - b. Drag a segment to the placeholder.

The placeholder is replaced by the segment.
 - c. To restrict access to additional segments, repeat these steps for each additional segment.
 - d. **Noprint Restriction.** To let the user access a field or segment for retrieval but not display its value on output, do the following:
 - a. Right-click an access right under the user you want to restrict, point to *Insert*, and select *Noprint Restriction*.

A Noprint restriction placeholder is added under the access right.
 - b. Drag a field to the placeholder.

The placeholder is replaced by the field.
 - c. To restrict print access to additional fields, repeat these steps for each additional field.
7. When you are finished adding restrictions to the synonym, close the DBA window and save the synonym.

In order to access the synonym, a user must issue a SET PASS command and enter one of the user passwords. This command can be issued in any supported profile or in a procedure.

Pivoting Columns

You can pivot columns that represent identical data or groups of such columns to display them as rows. This is useful when you want to access their values as one field. For example, if a spreadsheet has a column for each year, you may want to create a row containing all of the years for reporting.

How to Pivot Columns

The spreadsheet pivot_demo.xlsx is created as part of the WebFOCUS Retail tutorial and can be downloaded to a PC and then uploaded to demonstrate this feature.

1. Multi-select the columns by clicking in the data portion of the first column in the range you want to pivot, holding down the Shift key, and clicking in the data portion of the last column you want to pivot. Click the menu icon in one of the columns, and click *Multiple columns to rows*.

The Pivot dialog box opens, as shown in the following image.

Pivot Column(s) to Rows in synonym pivot_demo

Selected Parameters

? Pivot Type: Repeating column

? First column: 1960

? Last column: 2012

? Column Title for Pivoted Value: Pivot value (Example: AMOUNT)

? Column Title for Pivoted Key: Pivot key (Example: STMT_YEAR)

? Formula for Pivoted Key: 1960 + ORDER - 1

Cancel OK

2. Enter or select the following values.

Pivot Type. Select *Repeating column* or *Repeating group of columns*.

If you select *Repeating group of columns*, some of the information you need to enter changes.

- Number of groups.** This is the number of groups of columns for repeating groups of columns. For example, if an Excel spreadsheet has data that includes two columns of consecutive data for Plan and Actual figures for six years, this number would be 6.
- First Column.** This is the first column in a group of repeating columns.
- Last Column.** This is the last column in a group of repeating columns.
- First column in group.** Is the name of the first column in a group. For example, if an Excel spreadsheet has groups of data consisting of two columns called *year_plan* and *year_actual*, the first *year_plan* of the first group would be selected for *First column in group*.
- Last column in group.** Is the name of the last column in a group. For example, if an Excel spreadsheet has groups of data consisting of two columns called *year_plan* and *year_actual*, the first *year_actual* of the first group would be selected for *Last column in group*.
- Column Titles for Pivoted Value.** This is the name for the new column containing the data values. the default is *Pivot value*. For groups, separate the titles for the columns in the groups with slashes (/).
- Column Title for Pivoted Key.** This is the name of the new column that identifies the values in the data column. If the source is a spreadsheet this would be the name that could apply to the column headers. The default is *Pivot key*.
- Formula for Pivoted Key.** This is automatically calculated based on the values selected for First column/Last column or First column in group/Last column in group.

3. Click *OK*.

The data is pivoted using an occurs attribute. The following image shows the synonym and its output pane after pivoting the years from 1960 to 2012.

Country Name	Country Code	Indicator Name	Indicator Code
Afghanistan	AFG	1960 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1961 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1962 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1963 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1964 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1965 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1966 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1967 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1968 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1969 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1970 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1971 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1972 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1973 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1974 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1975 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1976 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1977 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1978 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1979 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1980 GDP per capita (current US\$)	NY.GDP.PCAP.CD
Afghanistan	AFG	1981 GDP per capita (current US\$)	NY.GDP.PCAP.CD

In the original synonym, each year was a column, but in the pivoted synonym, each year is a row.

Other Synonym Enhancements

Data Assist is an environment for managing and editing Synonyms. The number and types of changes you can make to synonyms has been greatly expanded.

Reference: Right-Click Options for Fields in the Table/Column, Business View, and Output Panes

You can rename a field by right-clicking the field and clicking *Rename* on the context menu. In the Business View frame, you can assign a DV role by right-clicking a field, pointing to *DV Role*, and clicking a role. Also in the Business View frame, you can right-click a field, point to *Manage Folders*, and click *New Folder* or *Create Parent Folder*.

You can also right-click a field to insert one of the following.


New Filter

You can open the Filters/Groups pane to create a filter for the field that will select only values satisfying the filter expression. For more information, see [Filters/Groups Frame](#) on page 287.

New Expression

- ❑ **Apply Function.** The types of functions available depends on the data type of the field. For more information about individual functions, see the *TIBCO WebFOCUS® Using Functions* manual or generate a Function Parameters report from the *FOCUS SETs and Info* item on the Workspace page of the Reporting Server browser interface.

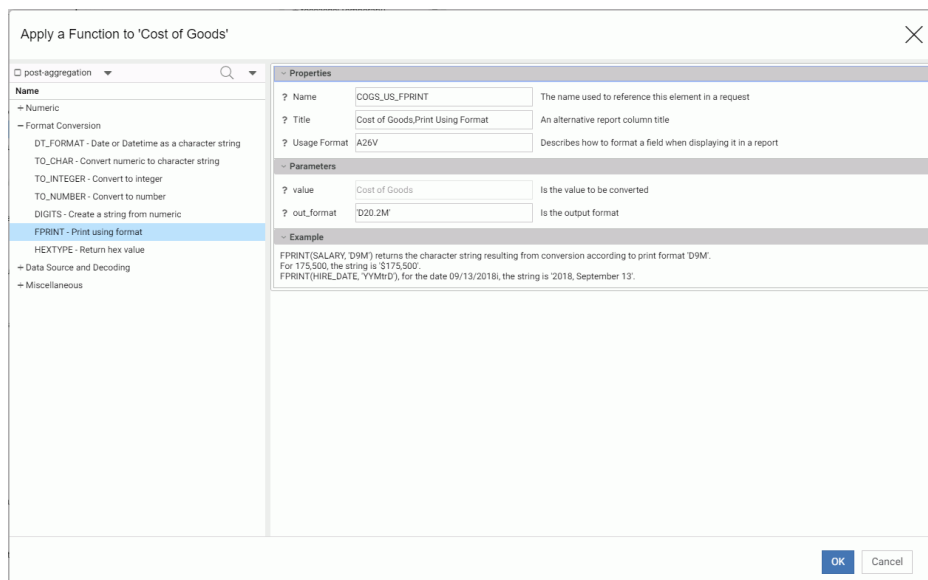
The Apply a Function dialog box opens, as shown in the following image.



The recommended categories of functions display in the left frame. By default, these are functions that apply to the selected field and you can use in a DEFINE virtual field. You can control which categories of functions are displayed using the post-aggregation check box to show functions that are applied after aggregation is performed. This will create a COMPUTE calculated value. The Advanced options menu (down arrow) also affects the categories of functions displayed. You can click the following options:

- Show legacy functions.** Adds legacy functions for the recommended categories to the left frame. If you click this option, the list item changes to *Hide legacy functions*.
- Show regardless of format compatibility.** Adds function categories that are not compatible with the format of the selected field. If you click this option, the list item changes to *Respect field type*.
- Hide function titles.** Shows only the function names on the list, not descriptive titles. If you click this option, the list item changes to *Show function titles*.

Click the function you want to apply. A dialog box opens in the right frame for entering or selecting function properties and parameters, as shown in the following image.



When you have entered the required information, click **OK**. The virtual field is added to the synonym.

Advanced Expression. Opens the expression calculator, described in [Using the Expression Calculator](#) on page 301.

- ❑ **Group and Replace.** Opens the Filters/Groups pane so that you can categorize the field values into groups for sorting. It creates a new virtual field using a conditional expression.

For more information about the Filters/Groups pane, see [Filters/Groups Frame](#) on page 287.


To open the Filters/Groups pane to edit the grouping, you can right-click the virtual field and select *Edit*.

To open the expression calculator for more advanced editing, right-click the virtual field and select *Edit (Advanced)*.

- ❑ **Phonetic Grouping.** This option is available for grouping alphanumeric fields based on one of two methods for encoding names based on pronunciation, Soundex or Metaphone. Soundex calculates a code for alphanumeric values so that if they sound the same but are spelled differently, they will still have the same code. Metaphone is an improved version of Soundex that takes into account anomalies in English spelling and pronunciation. During upload, this option can correct spelling inconsistencies before actually uploading the data.

For example, upload the file sales_orders_misspelled_cities.csv that is in the Uploads folder of the WebFOCUS Retail application.

Right-click the City field in the Geography dimension, point to *New Expression*, then *Groups*, and click *Phonetic Grouping*. The group will initially open in the Filters/Groups pane with the data values grouped based on the default phonetic algorithm for your operating environment, as shown in the following image.



The screenshot shows a table titled "City, Group" with two columns: "Groups/Values" and "Count". The data is grouped by city names, with each group containing the original name and its phonetic variants. The counts for each variant are listed in the "Count" column. The table is displayed in a scrollable pane with a search icon and a dropdown arrow at the top right. At the bottom, there are navigation controls including a minus sign, a plus sign, and "Page 1/3".

Groups/Values	Count
Albany	
Albany	109
Albani	1
Albeny	1
Beaumont	
Beaumont	290
Beaumunt	1
Barcelona	
Barcelona	496
Barcelone	2
Bartlesville	
Bartlesville	159
Bartlesvile	1
Decatur	
Decatur	369
Decater	1
Dayton	
Dayton	52
Daytan	1
Daytin	1
Fort Worth	
Fort Worth	8
FortWorth	1
Frankfort	
Frankfort	148
Frankfert	1
Frenkfort	1
Houston	
Houston	594
Hauston	1
Houstan	1
Jersey City	
Jersey City	524
Jersey Citi	2
Cleveland	
Cleveland	206
Clevelind	1

Each group contains different spellings of city names that would be pronounced the same way based on the phonetic algorithm used. If you do not agree with a grouping, you can edit it. Any data value that could not be placed in a group based on the phonetic algorithm is placed in a group called (Ungrouped Values). You can move those values into another group, if necessary.

A virtual field is created with an expression that causes all of the values within each group to be uploaded with the correct spelling, as shown in the following image.

Properties		
^ General		
? DEFINE	CITY_GROUP	An optional attribute used to create a virtual field for reporting
? EXPRESSION	IF City IN ('Albany', 'Albani', 'Albeny') THEN 'Albany' ELSE IF City IN ('B')	Is the calculation that creates the virtual field
? TITLE	City_Group	An alternative report column title for the field
? FORMAT	A19V	Describes how to format a virtual field when displaying it in a report or using it in a calculation
? Type	Character	Is the data type
? Length	19	Is a length specification
? MISSING	ON - missing only if all missable operands are missing	OFF - the result is never missing, missing value is replaced by 0 for numeric, blank for alphanumeric

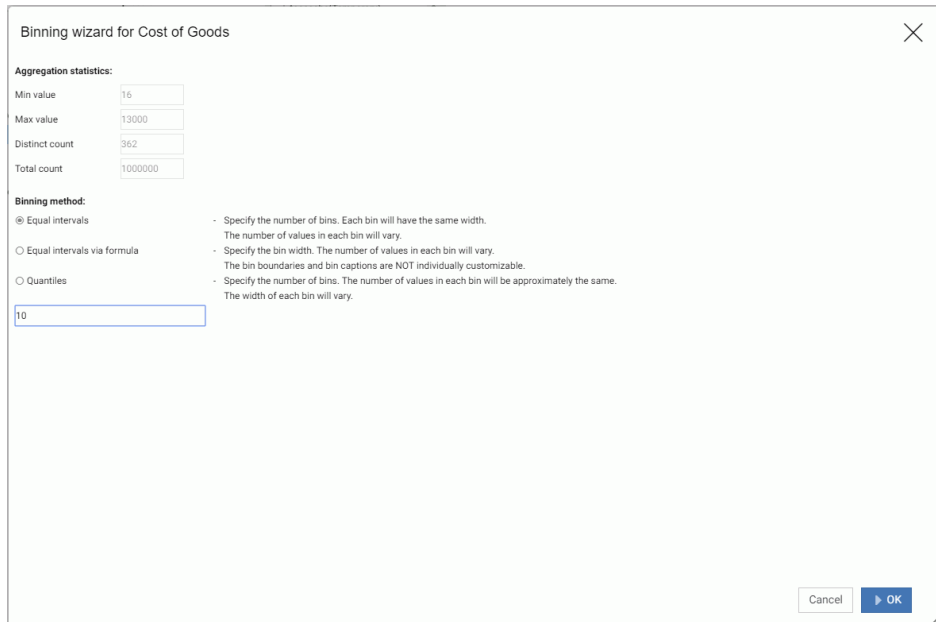
You can upload both the CITY and CITY_GROUP fields, or hide the CITY field and rename CITY_GROUP to CITY so that the CITY field has the corrected data values.

The options on the group menu specific to phonetic grouping are:

- Recreate groups by Metaphone.
- Recreate groups by SoundEx.

For information about other options for groups, see [Filters/Groups Frame](#) on page 287.

- ❑ **Create Bins.** Opens a wizard for dividing the data values in the field into a specified number of bins based on equal intervals, equal intervals using a formula, or quantiles, as shown in the following image.



The default number of bins is 10, but you can specify a different number for equal intervals and quantiles. For equal intervals via formula, you specify the bin width.

The virtual fields created using this option can be used for sorting.

The following image shows the binning card for bins based on equal intervals.

Discount, Bin (editable) <input checked="" type="checkbox"/> ▼	
High value	Bin caption
639	0639 & less
1278	0639 - 1278
1917	1278 - 1917
2556	1917 - 2556
3195	2556 - 3195
3834	3195 - 3834
4473	3834 - 4473
5112	4473 - 5112
5751	5112 - 5751
Max	5751 & more

Click *Apply* (check mark) to apply these bins.


The virtual field is added to the synonym.

The *Menu* (down arrow) has the following options:

- Add Bin.** This option adds another bin (the binning algorithm is recalculated using the new number of bins).
- Reset Bins.** This option opens the wizard with the original bin properties. Click *OK* to reset the bins, or edit the values and click *OK*.
- Edit properties.** Opens a dialog box in which you can edit the field name and field title of the generated virtual field.
- Hide Binning Card.** Closes the binning frame.

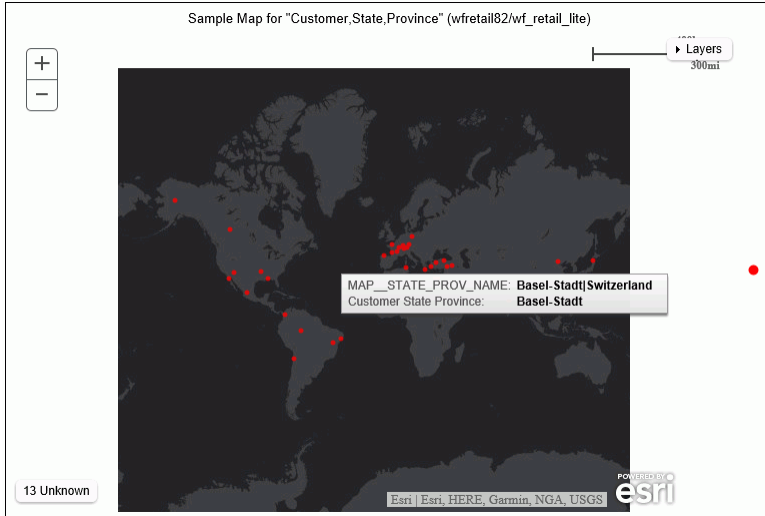
Geographic Role

You can right-click a field and assign it a geographic role, or change its geographic role assignment if it already has one, as shown in the following image.




To remove a geographic role, select (None).

For a field that has a geographic role assigned, you can generate a sample map by right-clicking the field and clicking *Sample Map* from the context menu. A sample map is shown in the following image.



Sample Data

You can right-click a field and select Sample Data to generate sample data for that field, as shown in the following image.



Partial Output is shown in the following image.

Segment	Name	Format	Count	Distinct Count	Distinct Percent	Minimum	Maximum	Average	Nulls Count	Nulls Percent	Duplicate Count	Duplicate Percent
MEASGRP1	REVENUE_US	D20.2M	10000	757	7.57	\$17.99	\$11,999.55	\$411.81	0	.00	9243	92.43

Data Profiling (Combo) and Data Profiling (Custom)

Data profiling provides reports and charts analyzing the values in a field or, for some reports, for a segment. For most reports, a *View* menu on the menu bar of the reports allows you to generate a printable view of the report, and a search function allows you to search for characters. Right-click a field or segment, point to *Data Profiling (Combo)* to get counts and a bar chart, or point to *Data Profiling (Custom)* to select one of the following options.

- Statistics.** This report lists counts, percents, limits, distinct values, and information about null values in a field or a segment, as shown in the following image.

Segment	Name	Format	Count	Distinct Count	Distinct Percent	Minimum	Maximum	Average	Nulls Count	Nulls Percent	Duplicate Count	Duplicate Percent
MEASGRP1	REVENUE_US	D20.2M	10000	757	7.57	\$17.99	\$11,999.55	\$411.81	0	.00	9243	92.43

Numeric columns are sorted numerically.

- Key Analysis.** This report lists counts and distinct and duplicate counts and percents for a field or a segment, as shown in the following image.

Segment	Name	Format	Elements	Count	Distinct Count	Distinct Percent	Duplicate Count	Duplicate Percent
MEASGRP1	QUANTITY_SOLD	I11C	1	10000	4	.04	9996	99.96
MEASGRP1	SALE_UNITY	I9C	1	10000	1	.01	9999	99.99
MEASGRP1	QUANTITY_SOLD/SALE_UNITY	I11C/I9C	2	10000	4	.04	9996	99.96

- Distinct Values.** Lists the 50 highest distinct values for numeric fields.

- Distinct Values Chart.** Shows a bar chart of distinct values for numeric fields.
- Frequent Values.** Lists frequent values for alphanumeric fields.
- Frequent Values Chart.** Shows a bar chart of frequent values for numeric fields.
- Frequent Values Pie Chart.** Shows a pie chart of frequent values for numeric fields.
- Distribution Chart.** Shows a bar chart of counts.
- Duplicate Values.** This report lists the values, count of duplicates, and corresponding percent, as shown in the following image.

Value	Count	Percent
\$17.99	1	.01
\$22.49	1	.01
\$23.99	5	.05
\$25.49	9	.09
\$26.99	9	.09
\$29.99	82	.82
\$30.00	1	.01
\$35.00	3	.03
\$37.50	1	.01
\$40.00	7	.07
\$41.99	1	.01
\$42.50	5	.05
\$45.00	17	.17

- Patterns.** This report shows the patterns of values for alphanumeric fields. The capital A character represents a capital letter, the lowercase a character represents a lowercase letter, and a 9 character represents a numeric character. Other characters represent themselves, as shown in the following image for the customer email field.

Pattern	Count	Percent
<u>AaaaaaAaaaa9999@aaaaa.aa</u>	105	1.05
<u>AaaaaAaaaa9999@aaaaa.aa</u>	101	1.01
<u>AaaaaAaaaa9999@aaaaa.aa</u>	100	1.00
<u>AaaaaAaaaa9999@aaaaa.aa</u>	97	.97
<u>AaaaaAaaaa9999@aaaaa.aa</u>	87	.87
<u>AaaaaAaaaa9999@aaaaa.aa</u>	83	.83
<u>AaaaaAaaaa9999@aaaaa.aa</u>	76	.76
<u>AaaaaAaaaa9999@aaaaa.aa</u>	68	.68
<u>AaaaaAaaaa9999@aaaaa.aa</u>	66	.66
<u>AaaaaAaaaa9999@aaaaa.aa</u>	66	.66
<u>AaaaaAaaaa9999@aaaaa.aa</u>	66	.66
<u>AaaaAaaaa9999@aaaaa.aa</u>	65	.65
<u>AaaaaAaaaa9999@aaaaa.aa</u>	65	.65
<u>AaaaAaaaa9999@aaaaa.aa</u>	64	.64
<u>AaaaaAaaaa9999@aaaaa.aa</u>	63	.63
<u>AaaaaAaaaa9999@aaaaa.aa</u>	62	.62
<u>AaaaaAaaaa9999@aaaaa.aaa</u>	61	.61
<u>AaaaaAaaaa9999@aaaaaaaa.aa</u>	61	.61
<u>AaaaaAaaaa9999@aaaaaa.aaa</u>	56	.56
<u>AaaaaAaaaa9999@aaaaaaaa.aa</u>	56	.56
<u>AaaaaAaaaa9999@aaaaa.aa</u>	55	.55

- Hex View.** This report lists the values in a field and their corresponding hexadecimal equivalents, as shown in the following image.

Value	Hex
AaliyahAkhtar4693@superrito.gb	41616C69796168416B68746172343639334073757065727269746F2E6762
AaliyahLittle3939@rhyta.sg	41616C697961684C6974746C65333933394072687974612E7367
AaliyahStone8968@superrito.gb	41616C6979616853746F6E65383936384073757065727269746F2E6762
AaronGregory2285@armyspy.sg	4161726F6E477265676F7279323238354061726D797370792E7367
AaronMarkley8756@armyspy.com	4161726F6E4D61726B6C6579383735364061726D797370792E636F6D
AaronShaffer7084@teleworm.us	4161726F6E53686166666572373038344074656C65776F726D2E7573
AbbeySparkes1594@superrito.sg	4162626579537061726B6573313539344073757065727269746F2E7367
AbbyDickinson3247@rhyta.gb	416262794469636B696E736F6E333234374072687974612E6762
AbdelhadivanGrinsven159@cuvox.nl	416264656C6861646976616E4772696E7376656E313539406375766F782E6E6C
Abdul-AzimWasem6328@rhyta.tr	416264756C2D417A696D576173656D363332384072687974612E7472
Abdul-QawiMoghadam2697@cuvox.tn	416264756C2D516177694D6F67686164616D32363937406375766F782E746E
Abdul-WahhabBatala1180@cuvox.tr	416264756C2D5761686861624261746131313830406375766F782E7472

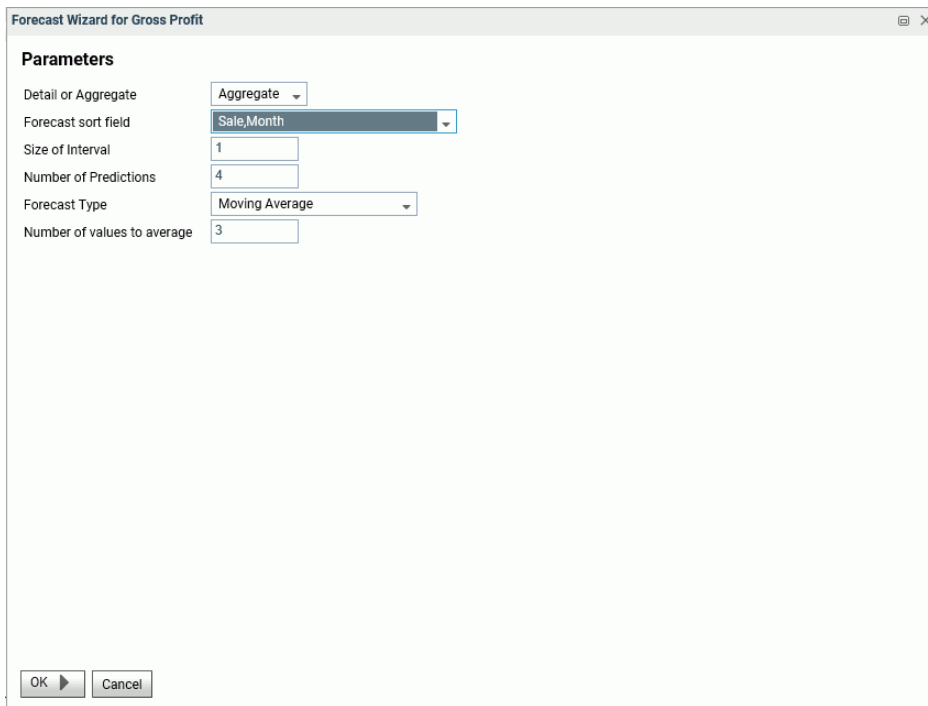
- Duplicate Values.** This report lists the values that have multiple instances, the count of instances, and the corresponding percent, as shown in the following image.

Value	Count	Percent
\$23.99	5	.05
\$25.49	9	.09
\$26.99	9	.09
\$29.99	82	.82
\$35.00	3	.03
\$40.00	7	.07
\$42.50	5	.05
\$45.00	17	.17
\$50.00	74	.74
\$50.98	2	.02
\$53.98	3	.03
\$59.49	2	.02
\$59.98	34	.34


- ❑ **Outliers.** This report shows high and low values with a count of instances that fall outside the bulk of the data for a field, as shown in the following image.

High	Count	Low	Count
\$11,999.55	1	\$17.99	1
\$9,447.30	1	\$22.49	1
\$7,999.70	2	\$23.99	5
\$7,998.00	2	\$25.49	9
\$6,998.00	2	\$26.99	9
\$6,798.00	6	\$29.99	82
\$5,438.40	1	\$30.00	1
\$4,499.98	1	\$35.00	3
\$4,199.85	1	\$37.50	1
\$4,198.80	1	\$40.00	7

- ❑ **Forecast chart.** Opens a Forecast dialog box for selecting parameters for the FORECAST functions, as shown in the following image.



You can choose to see a bar chart of moving averages or linear regression and enter a number of values to predict beyond the end of the data values, as shown in the following image.



Impact Analysis

Right-click a field and click *Impact Analysis* to generate a tab in the output frame that shows where and how the field is used, including the application name where the field is referenced, the file and type of file in which it is used, and the type of procedure, context, and line number in which it is found.

Key Analysis Using Sample Data

The Key Analysis report helps you select one or more fields that uniquely identify each record in a data set, so that these fields can be used as the key when loading the data to a relational database table.

Using Key Analysis, you select a synonym and then select fields that you think may be good candidates to become key fields. The server will then evaluate the fields, separately and in groups of two, three, and so on, to generate statistics that enable you to select an optimal set of key fields.

While using the entire data set ensures total accuracy in determining the uniqueness of the fields, for large data sets this could take a long time to process. Therefore, if you enable sampling in your Reporting Server browser interface preferences, a representative sample of the data will be used, enhancing performance while providing a reasonable assessment of the uniqueness of the selected fields.

After determining which fields should be key fields using the sample, rerun the analysis for those fields using the full set of data, in order to be sure that they uniquely identify each row.

The Key Analysis report is most useful when run against a data set that does not already have keys defined in the metadata. Most relational tables already have key definitions.

Procedure: How to Enable Sampling

1. Go to the Workspace page of the Reporting Server browser interface, click *Settings*, then click *Settings for Reporting Server browser interface Preferences*.

The *Change Settings for Reporting Server browser interface Preferences* page opens.

2. On the Workspace page of the Reporting Server browser interface, click *Settings*, then click *Settings for Web Console Preferences*.

The *Change Settings for Web Console Preferences* page opens.

3. Expand the *Data Assist (Representative Sampling)* section as shown in the following image.

Data Assist (Representative Sampling)

ENABLE_SAMPLING ?

 Controls Representative Sampling in Data Flow

SAMPLING_APPROOT ?

 Application for representative samples. Default: foccache

SHOW_SAMPLING_INFO ?

 Show Sampling Info

SMPL_ETL_TRG_DBMS ?

 Use ETL-TRG-DBMS for Sampling

SMPL_PRIORITIZE_JOIN ?

 Prioritize join when creating sampling

SAMPLING_SEED ?

 Seed for representative samples

SAMPLING_TARGET_COUNT ?

 Maximum record count for representative samples

4. Select *On* for *ENABLE_SAMPLING*.

This enables sampling everywhere it is available.

You can also change the following parameters. Not all of them are related to Key Analysis.

SAMPLING_APPROOT

Is the application where the sampling files will be generated. The default is foccache.

SHOW_SAMPLING_INFO

This shows sampling information on the data flow canvas. The default value is Yes.

SMPL_TRG_DBMS

Use the adapter set as ETL_TRG_DBMS for sampling. The default value is No.

SMPL_PRIORITIZE_JOIN

Prioritize the join when creating sampling. When this is set to Yes, the sampling for a dimension table in a join will contain all of the key values from the sample rows generated for the fact table to which it is joined. The default value is No.

SAMPLING_SEED

It is used to get repeatable representative sampling results. The default value is empty, meaning no seed is used and each time sampling is re-created it will have different data.

SAMPLING_TARGET_COUNT

It shows maximum record count for representative samples. This field is empty by default. The value of this field is calculated automatically.

5. Click Save.

Procedure: How to Run the Key Analysis Report

The samples shown in this section use a .csv file uploaded from the Citibike bike share web site that contains data about bike rides during a specific month.

1. Right-click a synonym, point to *Data Profiling*, and click *Key Analysis*.

The Key Analysis page opens, as shown in the following image.

<input type="checkbox"/>	Segment Name	Display Field Name	Data Type	Nulls	Description
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	tripduration	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	start station id	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	end station id	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	bikeid	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	birth year	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	gender	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime	Date and Time	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	stoptime	Date and Time	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	usertype	Character	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Year	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Quarter	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Month	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Day	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y	Date (Y)	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y-Q	Date (YQ)	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y-M	Date (YM)	Yes	

Analyze (All Data) Analyze

2. Select the check boxes next to the columns you want to have analyzed, and click *Analyze*.

The key analysis report opens, showing statistics for all combinations of the columns you selected, as shown in the following image.

Key Analysis for "start station id"/"end station id"/bikeid/starttime/stoptime/usertype (CITIBIKE/CITIBIKE_TRIPDATA_CSV)

Segment	Name	Data Type	Elements	Count	Distinct Count	Distinct Percent	Duplicate Count	Duplicate Percent
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id</u>	Integer	1	16466	777	4.72	15689	95.28
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id</u>	Integer	1	16466	777	4.72	15689	95.28
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid</u>	Integer	1	16466	9096	55.24	7370	44.76
CITIBIKE_TRIPDATA_CSVOUT	<u>starttime</u>	Date and Time	1	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>stoptime</u>	Date and Time	1	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>usertype</u>	Character	1	16466	2	.01	16464	99.99
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/end station id</u>	Integer/Integer	2	16466	13715	83.29	2751	16.71
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/bikeid</u>	Integer/Integer	2	16466	16419	99.71	47	.29
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/usertype</u>	Integer/Character	2	16466	1421	8.63	15045	91.37
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/bikeid</u>	Integer/Integer	2	16466	16432	99.79	34	.21
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/usertype</u>	Integer/Character	2	16466	1403	8.52	15063	91.48
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00

Clicking a column heading sorts on that column, toggling between ascending and descending order.

The following columns are included in the report.

Segment

Specifies the segment in the synonym that contains the field.

Name

Is the field name or title of the field or combination of fields, depending on your name display strategy.

You can right-click a row in the Name column and select one of the following:

- Duplicate values.** Shows the count and percent of duplicate values for each combination of field values.
- Duplicate values chart.** Available for single fields only. Displays a horizontal bar chart showing the count of duplicate values for each field value.
- Distribution chart.** Available for single fields only. Displays a vertical bar chart showing the count of values for each range of field values.
- Duplicate values pie chart.** Available for single fields only. Displays a pie chart showing the count of duplicate values for each field value.

Data Type

Is the data type or combination of data types for the fields listed on each row.

Elements

Is the number of fields included in the analysis for each row.

Count

Is the number of instances of the field or combination of fields for each row, based on the sample.

Distinct Count

Is the number of distinct instances of the field or combination of fields for each row, based on the sample.

Distinct Percent

Is the percent of distinct instances of the field or combination of fields for each row, based on the sample.

Duplicate Count

Is the number of duplicate instances of the field or combination of fields for each row, based on the sample.

You can right-click a row in the Duplicate Count column and select one of the following:

- Duplicate values.** Shows the count and percent of duplicate values for each combination of field values.
- Duplicate values chart.** Available for single fields only. Displays a horizontal bar chart showing the count of duplicate values for each field value.
- Distribution chart.** Available for single fields only. Displays a vertical bar chart showing the count of values for each range of field values.
- Duplicate values pie chart.** Available for single fields only. Displays a pie chart showing the count of duplicate values for each field value.

Duplicate Percent

Is the percent of duplicate instances of the field or combination of fields for each row, based on the sample.

3. Look for fields or combinations of fields on the report that provide 100% distinct values.

In order to be a key, the combination of fields must provide 100% distinct values. Select the smallest combination of fields that provide distinct values and that are good candidates to be key fields.

- ❑ If no combination of fields provides 100% distinct values, either add additional fields to the Key Analysis or create a field (such as a row number) that will provide a distinct value for each row.
- ❑ If multiple fields or combinations of fields on the report provide 100% distinct values, use your knowledge of the data to select a good combination.

For example, in the sample report shown, both the starttime and endtime fields provide 100% distinct values, but that may be an anomaly due to the sample chosen. There is always the possibility that multiple bike rides will start or end at the same time.

Similarly, the combination of starttime and endtime may not always provide unique values. The combination of bikeid with either starttime or endtime is probably unique, as the same bike could not be used for multiple rides at the same exact date and time. Also, combinations that include more than two fields are not as desirable as combinations that include only two fields.

4. Once you have chosen a key, rerun the Key Analysis report, selecting only those candidate fields, and this time clicking *Analyze (All Data)*.

Check to make sure that this selection provides a unique key for all of the data, not just the sample. In this example, the combination bikeid/starttime provides a unique key, as shown in the following image.

Key Analysis for bikeid/starttime (CITIBIKE/CITIBIKE_TRIPDATA_CSV)

Segment	Name	Data Type	Elements	Count	Distinct Count	Distinct Percent	Duplicate Count	Duplicate Percent
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid</u>	Integer	1	2181064	14356	.66	2166708	99.34
CITIBIKE_TRIPDATA_CSVOUT	<u>starttime</u>	Date and Time	1	2181064	2179730	99.94	1334	.06
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/starttime</u>	Integer/Date and Time	2	2181064	2181064	100.00	0	.00

Working With Data Flows

A data flow contains the parameters for extracting information from data sources and loading it into data targets. You create data flows in the Reporting Server browser interface using the data flow workspace. The data flow generates an SQL request that extracts data from the sources, transforms it to meet the specifications of the target, and loads it into the target.

Typically, you build a data flow from a set of objects that identify data sources, data targets, data columns to include, and business rules governing how data is extracted, transformed, and loaded to form a data warehouse, data mart, or any other type of target.

You can join or union multiple data sources, filter and sort the extracted source data, and specify column mappings from the source to the target.

In this chapter:

- [Creating a Data Flow](#)
 - [Preparing Data With a Flow](#)
 - [Editing Data Flow Options](#)
 - [Assigning Slowly Changing Dimension \(SCD\) Columns](#)
 - [Generating Data Flows Using a Template Flow](#)
 - [Controlling Format Conversion Error Processing in a Data Flow](#)
 - [Migrating Data Flows](#)
 - [Viewing Checkpoint Files Generated by a Change Data Capture \(CDC\) Process](#)
-

Creating a Data Flow

The first step in building a data flow is to create an empty flow. You can also create a flow by right-clicking a synonym and clicking *Flow*, which is easier because the source is pre-selected. You will then be able to add objects to this flow, and set up mappings and relationships.

Before loading data from the source to the target, you can prepare it for loading by implementing enhancements and transformations, as described in [Preparing Data With a Flow](#) on page 376.

Procedure: How to Create a Data Flow

You can build data flows ranging from the simple, where data is moved from one source into one target, to the complex, which employs multiple sources and targets and a variety of data extraction techniques.

You can view the generated SQL for a flow in a pop-up window by clicking the *Advanced Options* menu, clicking *Diagnostics*, and selecting *View Source*).

1. Right-click the application directory where you want the new data flow.
2. Point to *New* and then click *Flow*.


or

Right-click a synonym and click *Flow* to open a new flow with this synonym inserted as a source.

or

Click the *New* button on the Applications page ribbon, and click *Flow*.

An empty data flow opens in the data flow workspace, with an SQL object in the center of the workspace, as shown in the following image.



The steps to creating and running a data flow are:

1. Add the first source to the flow and select the columns from the source that you want to extract and load to the target.
2. Optionally, add additional sources.

Optionally, prepare the data for loading to the target by enhancing the data with calculations.

3. Add one or more targets to the flow.

4. Optionally, edit the load options for the flow.
5. If you are merging new records into an existing target, optionally edit the merge conditions for the flow.
6. Run the flow.

Procedure: How to Add the first Source Object to the Data Flow

To add a source to the flow, select Data on the sidebar to show the available folders or synonyms.

1. To add the first source object to the flow, do one of the following.
 - Drag a data source synonym from the data pane into the data flow workspace.
or
 - Right-click the SQL object, point to *Add*, and click *Source*, *Join*, *Select*, or *Union*.
or
 - Right-click a *Join*, *SQL*, or *Union* object and click *Add*. Select the object from the available list.

Sources are represented by database-specific icons so that you can easily identify the type of source.

By default, all columns in a single-segment source or all columns in the top segment of a multi-segment source are added to the flow because *Automatically select all columns* is selected in the Advanced Options dialog box. You can change this so that no columns are automatically added by deselecting this option.

2. If you want to edit the list of columns that were automatically selected, or change their roles in the generated query, right-click the SQL object and click *Edit* to select the columns that you want to extract from the source and copy to the target.
3. Drag fields from the Segments/Fields pane to the Columns/Filters pane.

The selections will create a query for loading the target data source. You can drag fields to the Order By, Columns, Order Across, and Filters and Variables text boxes.

The Order By text box has a menu that has the option *Switch to Group By*, which generates an aggregated request instead of a detailed request. If you switched to Group By, the menu option becomes *Switch to Order By*.

The Columns text box has a menu with the option *Apply Function*, which opens the Function dialog box.

Reference: Distinguishing Node Status Using Ports and Connectors

The Data Flow and Join nodes have ports that indicate if they can have inputs (on the left) or outputs (on the right). The ports also indicate if all required inputs and outputs are configured.


- ❑ An empty node that is not yet considered missing is represented by a hollow circle.
- ❑ A node that is present is represented by a solid blue circle.
- ❑ A required node that is missing is represented by a circle half-filled in red and outlined in red.
- ❑ A node with a menu has an ellipsis (...) at the bottom right. You can click the ellipsis or right-click anywhere in the node to open the menu.
- ❑ The indicator that there are warning messages is an exclamation point in a red circle, which you can hover over to display the warning messages.

In the following image of a new flow, a source is required, but not present, indicated by the half-filled red circle on the left of the SQL node. No target has been added, so the circle on the right of the SQL node is hollow.



Select Columns

In the following image, a source has been added, so the circles between the source and the SQL node are solid blue. No target has been added, so the circle on the right of the SQL node is hollow.




citibike_tripdata_mssql(T1)

Select Columns

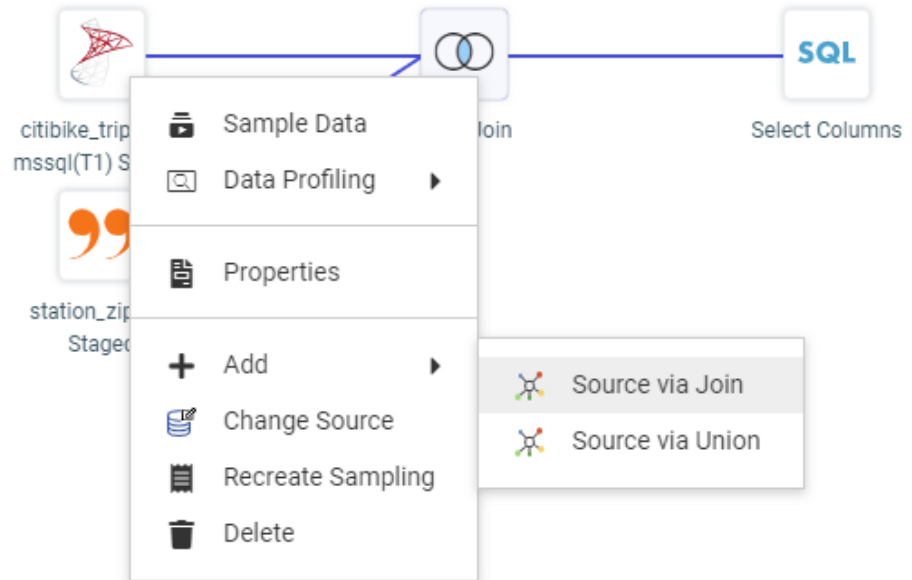
Procedure: How to Add Additional Sources to the Flow (Optional)

When you add additional sources to a flow they can be connected to an existing source using a join or union. A join adds columns to existing rows based on a relationship between columns in the two data sources, while a union adds new rows by combining the results of two queries that have a common domain and number of columns.

1. Perform one of the following actions to add a new source.
 - Drag a synonym from the data pane to the Join or Union bubble of the synonym to which you want to connect it, as shown in the following image.



- ❑ Right-click an existing source, click *Add* from the context menu, then *Source via Join* or *Source via Union*, as shown in the following image.



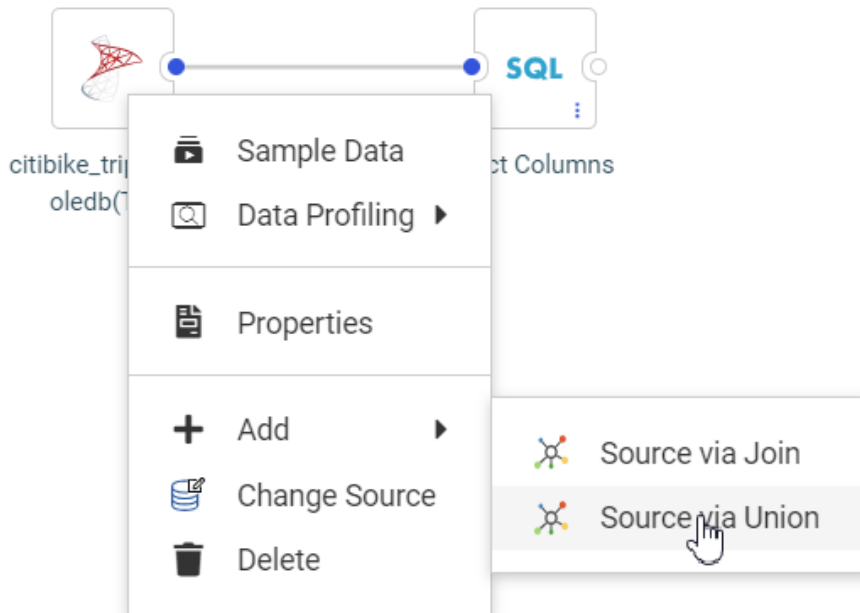
2. Optionally, edit the generated join or union.

To edit a join, right-click the join object, and click *Join Editor* from the context menu. The Join Editor is described in [Creating a Join](#) on page 308.

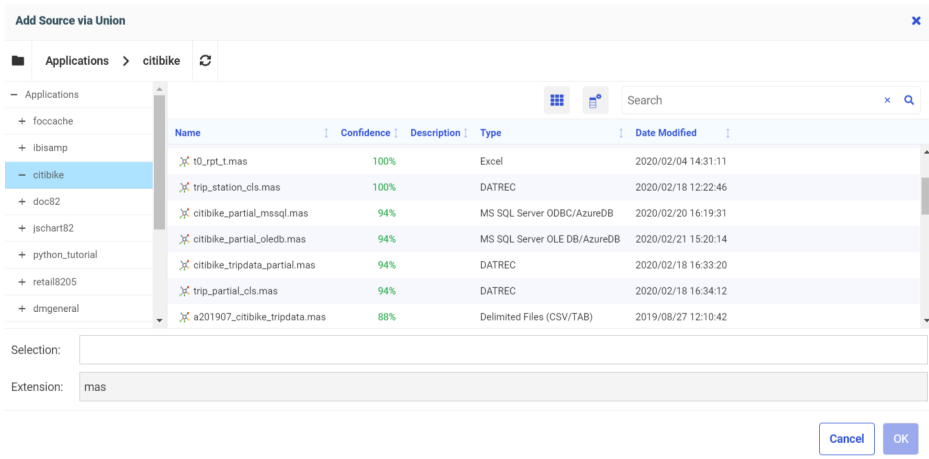
To edit a union, right-click the union object, and click *Union Editor* from the context menu.

Showing the Confidence Level for Adding a Union to a Data Flow

You can right-click a source in a data flow and add a new source as a union, as shown in the following image.



A list of tables displays for choosing the new source to add. This table now includes a confidence level, based on matching column names and data types for each prospective new source, showing how compatible it is with the existing source, as shown in the following image.



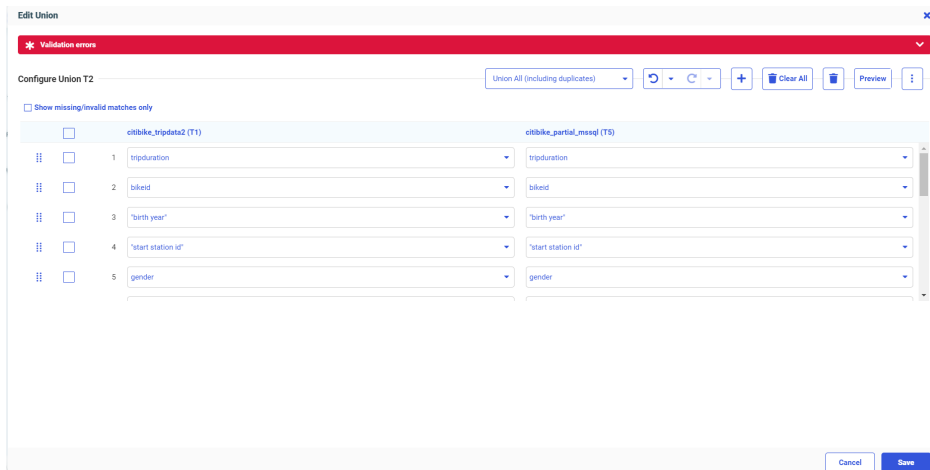
Using the Union Editor

When you right-click a Union object and select *Union Editor* from the context menu, the *Edit Union* page opens, showing the matching columns and any validation issues.

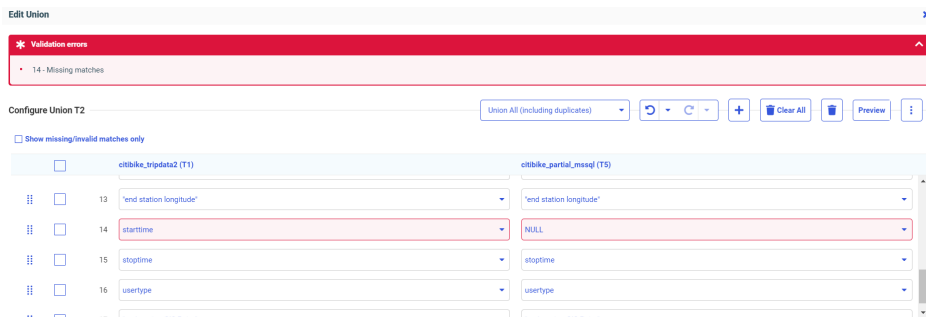
Reference: Correcting Validation Errors and Editing the Matches

The Union Editor has two areas. The top area shows the matches between the fields in the two tables. The bottom area contains a preview of the data in the union. When you first open the Union Editor, the Preview area is hidden.

If any validation errors exist, a message appears at the top of the window, as shown in the following image.



You can click the down arrow in the message area to see a list of problems. In the following image, line 14 has a missing match and, when you scroll down to that line, it is outlined in red.



You can also select the check box *Show missing/invalid matches only*, to display only those lines with missing or invalid matches.

To correct a missing match, you can:

- Click the down arrow next to the right table and select a different field that matches the field from the left table.
- Click the down arrow next to the left table and select a different field that matches the field from the right table.
- Select the check box next to the line and click the *Delete selected columns* button (trash can) at the top of the page.

You can select multiple columns and delete them all at once.

- Click the *Clear All* button (trash can with the text *Clear All*) at the top of the page to remove all columns.

You will then need to add pairs manually using the *Add a pair (+)* button.

- Click the *Add a pair (+)* button to add a matching pair.


Clicking this button adds a line with NULL selected for each table. Select fields from the left and right tables using the drop-down lists, to create a matching pair.

- Change the type of union generated, if that is appropriate. For information, see [Selecting the Type of Union](#) on page 355.

When all validation problems are resolved, the message area disappears.

Reference: Selecting the Type of Union

By default, Union All, including duplicates, is selected. However, you can change the type of union that is generated by selecting an option from the drop-down list at the top of the page, as shown in the following image.








A union combines the output of two SELECT statements into one answer set, where the number of columns in each is the same, and the data types are compatible. You can select the following types of union:


- Union (excluding duplicates).** Returns the rows from both SELECT statements with duplicate rows removed.
- Union All (including duplicates).** Returns all rows from both SELECT statements, including duplicates. This is the default union type.
- Except (excluding duplicates).** Returns distinct rows from the left SELECT statement that are not returned by the right SELECT statement.
- Except All (including duplicates).** Returns all rows from the left SELECT statement that are not returned by the right SELECT statement.

- Intersect (excluding duplicates).** Returns matching rows from the left SELECT statement and the right SELECT statement based on matching all columns. Duplicate rows are removed.
- Intersect All (including duplicates).** Returns matching rows from the left SELECT statement and the right SELECT statement based on matching all columns. Duplicate rows are retained.

Reference: Editing the Union Properties and View

The following table list changes you can make to the union properties and view.

Operation	Icon	Description
Move columns		Click to drag a pair up or down on the list.
Undo		Undoes the most recent change made to the union properties. You can click the down arrow next to the button to see a list of the changes that have been made, and select from the list. Selecting an item undoes that change and all prior changes.
Redo		Redoes the most recent change made to the union properties. You can click the down arrow next to the button to see a list of the changes that have been made and, select from the list. Selecting an item redoes that change and all prior changes.
View		Provides the following options. <ul style="list-style-type: none"> <input type="checkbox"/> Show title. Displays the field title. This is the default view. <input type="checkbox"/> Show Name. Displays the field name. <input type="checkbox"/> Show source name. Displays the name from the source file. <input type="checkbox"/> Diagnostics. Has two options, Session Log, which opens the Session Log, and Clear Session Log, which clears the Session Log.
Choose Columns		In the Preview area, this enables you to select the columns that display.

Operation	Icon	Description
Sort		In the Preview area, for each column, changes the sort order of the rows based on the data in that column. The sort toggles between the original order, descending, and ascending.

Reference: Viewing Sample Data

Click the *Preview* button at the top of the page to display sample data for the union, as shown in the following image.

The *Preview* button changes to a *Hide* button, which you can click to hide the sample data.

You can enter characters in the *Search* box to display lines with data values that match those characters.

You can use the *Choose Columns* button to select which columns display in the preview.

You can use the *Sort* button in a column to sort the rows based on that column.

Reference: Using Metadata Classifications in a Union


Prior to using this feature, the classification APIs must have been installed and configured, as described in your *TIBCO WebFOCUS® Installation and Configuration* manual. If you uploaded the files to be used in the union and clicked the *Classify* button, the classifications will have been added to the synonym.

We will upload two Excel files, and examine the classifications, then use the two Excel files in a Union.


In the server environment, you can upload a file by clicking *Get Data* on the Applications page of the Reporting Server browser interface.

Classifying Excel1

Click *Excel* in the Local Files section, and a file picker dialog box opens, as shown in the following image.

File name: Custom Files (*.xls;*.xlsx;*.xlsm)*.

Selecting *Excel1* and clicking *Open* opens the Upload Data page. If the Metadata Classification feature has been installed and configured, the Upload Data page has a *Classify* button, as shown in the following image.

CUSTOMER (PREVIEW/MODIFY TABLE)
Crosstab Formatted Raw Skipped rows: 0 Header rows: 1 More Search
ID Customer | Email | Address | First Name | Full Name | Name | G | Customer Address Line 1 | Customer City | Customer State Province | Customer Postal Code | Customer Country
167 ConnorHammond4443@eirrot.gb | Connor | Connor Hammond | Hammond | M | Computer Block | Port Talbot | England | SA13 2NG | United Kingdom
1800 MonicaWeiss6693@superrito.com | Monica | Monica Weiss | Weiss | F | 10 Rue De Bromont | Gatineau | Quebec | J8T 6E8 | Canada
3278 LisaRobinson7184@jourapide.com | Lisa | Lisa Robinson | Robinson | F | 680 E. Swedesford Rd. | Wayne | Pennsylvania | 19087 | United States
4247 StarrSmith9670@qustr.com | Starr | Starr Smith | Smith | F | 2843 Spring Lks | Missouri City | Texas | 77459 | United States
6725 KatherineHinojos1603@rhyta.com | Katherine | Katherine Hinojos | Hinojos | F | 147 Afton Ave | Youngstown | Ohio | 44512 | United States

Clicking the *Classify* button opens the Recommended Classifications page, as shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer			167.374849..(121)		
Email_Address		None Found	AaronConway8532@teleworm.ca, WolfgangBumgarner4230@enrot.de..(121)	AaAaaa9999@aaaaaa.aaa, AaaaaaaaAaaaaaa999@aaaaaaa.aa..(98)	W G W W
FirstName	85%	FIRSTNAME	Aaron, Wolfgang..(117)	Aa, Aaaaaaaa..(6)	W
FullName		None Found	Aaron Conway, Wolfgang Bumgarner..(120)	Aa Aaaaa, Aaaaaaaa Aaaaaaaa..(42)	W, W, W, W, W, W
Name		None Found	Allen, Winther..(116)	A Aaaa, Aaaaaaaa..(11)	W, W, W, W
G	88%	GENDER	F, M	A	W
CustomerAddress_Line 1	54%	ADDRESS_1	1 Queen's Crescent, Zitone Airport..(121)	9 A 99Aa Aa, Aaaaaaaa..(112)	1 1 W, W, W, W, W
CustomerCity	76%	ADDRESS_CITY	Abbeyleix, Zwolle..(109)	AA-Aaaaaa, Aaaaaaaa..(27)	W, W, W, W..(5)
CustomerState_Province	83%	ADDRESS_STATE	Alaska, Zurich..(71)	AaAaaa, Aaaaaaaa..(30)	W, W, W, W..(4)
CustomerPostal_Code		None Found	..W1F 0UJ..(66)	..AA99 9AA..(7)	..WWW W ..(7)
CustomerCountry	85%	ADDRESS_COUNTRY	Austria, United States..(18)	Aaaa, Aaaaaaaa..(6)	W, W, W

Note that not all columns were classified. You can select from the list of classifications for a column, or enter your own classification, by typing it into the text box next to the column. The columns that were classified have a Confidence Level associated with them. The confidence level represents the probability that the classification is correct.

In this case, select FULLNAME for Full,Name and LASTNAME for Name, as shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer			167.374849..(121)		
Email_Address		None Found	AaronConway8532@teleworm.ca, WolfgangBumgarner4230@enrot.de..(121)	AaAaaa9999@aaaaaa.aaa, AaaaaaaaAaaaaaa999@aaaaaaa.aa..(98)	W G W W
FirstName	85%	FIRSTNAME	Aaron, Wolfgang..(117)	Aa, Aaaaaaaa..(6)	W
FullName		FULLNAME	Aaron Conway, Wolfgang Bumgarner..(120)	Aa Aaaa, Aaaaaaaa Aaaaaaaa..(42)	W, W, W, W, W, W
Name		LASTNAME	Allen, Winther..(116)	A Aaaa, Aaaaaaaa..(11)	W, W, W, W
G	88%	GENDER	F, M	A	W
CustomerAddress_Line 1	54%	ADDRESS_1	1 Queen's Crescent, Zitone Airport..(121)	9 A 99Aa Aa, Aaaaaaaa..(112)	1 1 W, W, W, W, W
CustomerCity	76%	ADDRESS_CITY	Abbeyleix, Zwolle..(109)	AA-Aaaaaa, Aaaaaaaa..(27)	W, W, W, W..(5)
CustomerState_Province	83%	ADDRESS_STATE	Alaska, Zurich..(71)	AaAaaa, Aaaaaaaa..(30)	W, W, W, W..(4)
CustomerPostal_Code		None Found	..W1F 0UJ..(66)	..AA99 9AA..(7)	..WWW W ..(7)
CustomerCountry	85%	ADDRESS_COUNTRY	Austria, United States..(18)	Aaaa, Aaaaaaaa..(6)	W, W, W

Click *Apply Changes*. A green message bar at the top of the page shows that the changes were successfully applied.

Click *Load*.

The file is uploaded, and the synonym is generated with the classifications added, as shown in the following partial synonym.

```
FILENAME=EXCEL1, SUFFIX=MSODBC , BV_NAMESPACE=OFF, $
SEGMENT=EXCEL1, SEGTYPE=S0, $
  FIELDNAME=ID_CUSTOMER, ALIAS=ID_CUSTOMER, USAGE=I9, ACTUAL=I4,
  MISSING=ON,
  TITLE='ID Customer', $
  FIELDNAME=EMAIL__ADDRESS, ALIAS=EMAIL__ADDRESS, USAGE=A42V,
ACTUAL=A42V,
  MISSING=ON,
  TITLE='Email ,Address', $
  FIELDNAME=FIRST_NAME, ALIAS=FIRST_NAME, USAGE=A11V, ACTUAL=A11V,
  MISSING=ON,
  TITLE='First,Name',
  CATEGORY='FIRSTNAME', CONFIDENCE=0.853071, $
  FIELDNAME=FULL_NAME, ALIAS=FULL_NAME, USAGE=A26V, ACTUAL=A26V,
  MISSING=ON,
  TITLE='Full,Name',
  CATEGORY='FULLNAME', $
  FIELDNAME=NAME, ALIAS=NAME, USAGE=A15V, ACTUAL=A15V,
  MISSING=ON,
  TITLE='Name',
  CATEGORY='LASTNAME', $
  FIELDNAME=G, ALIAS=G, USAGE=A1V, ACTUAL=A1V,
  MISSING=ON,
  TITLE='G',
  CATEGORY='GENDER', CONFIDENCE=0.883651, $
  FIELDNAME=CUSTOMER_ADDRESS__LINE_1, ALIAS=CUSTOMER_ADDRESS__LINE_1,
USAGE=A57V, ACTUAL=A57V,
  MISSING=ON,
  TITLE='Customer,Address ,Line 1',
  GEOGRAPHIC_ROLE=ADDRESS_LINE,
  CATEGORY='ADDRESS_1', CONFIDENCE=0.547940, $
```

Classifying Excel2

Follow the same steps for file Excel2. The Classification Screen is shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer		None Found	167.374849..(121)		
Email Address		None Found	AaronConway8532@teleworm.ca,WolfgangBumgarner4230@ehnot.de..(121)	AaAaaa9999@aaaaaa.aaa,AaaaaaaaAaaaaaa999@aaaaaaa.aa..(98)	W@W W
Nme	85%	FIRSTNAME	Aaron,Wolfgang..(117)	Aa,Aaaaaaaa..(6)	W
FullName		None Found	Aaron Conway,Wolfgang Bumgarner..(120)	AaAaaaa,AaaaaaaaAaaaaaa..(42)	W,W,W W W
LastName		None Found	Allen,Wither..(116)	AAaaaa,Aaaaaaaa..(11)	W,W W W
Gender	88%	GENDER	F,M	A	W
CustomerAddress_Line 1	54%	ADDRESS_1	1 Queen's Crescent,Zitone Airport..(121)	9 A 95Aa Aa,Aaaaaaaa..(112)	W W W W,W W
CustomerCity	76%	ADDRESS_CITY	Abbeyleix,Zwolle..(109)	AA-Aaaaaaa,Aaaaaaaa..(27)	W,W W W..(5)
CustomerState_Province	83%	ADDRESS_STATE	Alaska,Zurich..(71)	AaAaaaa,Aaaaaaaa..(30)	W,W W W..(4)
CustomerPostal_Code		None Found	..W1F 00U..(66)	..AA99 9AA..(7)	..W1W (W)..(7)
CustomerCountry	85%	ADDRESS_COUNTRY	Austria,United States..(18)	Aaaaa,Aaaaaaaa..(6)	W,W W W

Note that the first name column is called Nme in this file (it was called Name in Excel1), and the gender column is called GENDER (it was called G in Excel1).

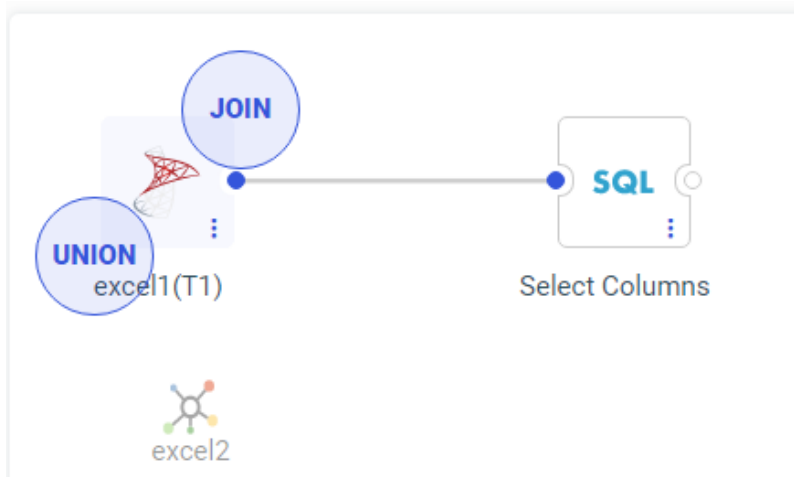
Again, select FULLNAME for Full,Name and LASTNAME for Name, click *Apply Changes*, and then *Load*.

Creating a Union

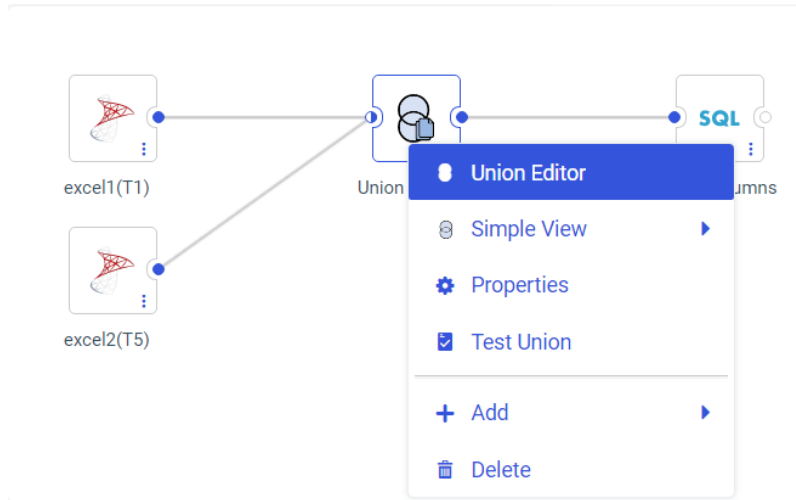
Now that the files have been uploaded and the synonyms have been generated with the classifications, click + *Data* on the Reporting Server browser interface ribbon, and then click *Flow*.

The Data Flow canvas opens.


Drag Excel1 to the flow canvas, then drag Excel2 to the canvas and select Union, as shown in the following image.



A Union object is added to the flow. Right-click the Union object, and click *Union Editor*, as shown in the following image.



The Union Editor opens, as shown in the following image.



The classified fields are correctly matched. Fields that were not classified, such as Customer ID and Email Address, are matched based on column names and formats. If any matches are missing or incorrect, you can correct them using the drop-down lists.

Reference: Saving the Union

When you have resolved all validation errors and made any edits you need, click *Save*.

Click *Cancel* if you do not want to save your changes.


Procedure: How to Add One or More Targets to the Flow

On the Data Flow sidebar, click *Targets*, then select *New Target* or *Existing Target*.


You can also drag targets onto the flow canvas:

- For a new target, drag the *New Target* icon onto the Data Flow canvas.
- There are two ways of dragging an existing target. You can drag the *Existing Target* icon from the *Targets* palette onto the canvas, which then opens a dialog box where you can select the target. You can also drag a synonym from the *Resources* pane to the right side of the SQL object.

If you select *New Target*, a Target object is added to the flow with a default name, as shown in the following image.




If you select the *Existing Target* icon, a file picker opens for selecting the existing target. The file picker shows the confidence level, based on matching column names and data types for each prospective new target, showing how compatible it is with the existing source, as shown in the following image.



The following options and icons apply to targets.

- You can rename the target by right-clicking it and clicking *Load Options*.
- A synonym will be generated for the target.
- Targets are represented by database-specific icons.
- For a new target, a plus sign (+) displays in the node.
- If the target has one or more keys, a key icon displays in the node, as shown in the following image of a new target with a key.



Before loading the data, you can edit the load options, as described in [How to Edit the Load Options for the Target \(Optional\)](#) on page 364.

If you are merging new records into an existing target, you can edit the merge conditions, as described in [How to Edit the Merge Conditions for an Existing Target \(Optional\)](#) on page 371.

Procedure: How to Edit the Load Options for the Target (Optional)

You can edit the options for loading the target or use the default options.

1. Right-click the target for which you want to edit the load options, and click *Load Options*.

The Load Options dialog box opens, as shown in the following image.

If any validation errors occur, the error messages display in red at the top of the screen. When you correct the problems, the messages disappear.

2. You can edit the following options.

Load Option

Choices are:

- New/Replace.** Create a new target or replace the target, if it exists. This is the default load option.
- New/Replace via SQL SCRIPT.** Instead of loading the data to a table, this option generates a SELECT statement and metadata that describes it. This option requires all the sources to be in the same database and use the same adapter and connection. Also, the query must be able to be completely optimized, that is passed through to the database for processing. This view can then be used as a source in a subsequent flow so that data is only retrieved once, without creating intermediate tables.
- Append to Existing.** Add new records after all existing records.

- Merge into Existing.** Merge the new records with the existing records. This option is not supported for delimited, DATREC, or JSON targets. Support with other types of targets depends on the type of write access WebFOCUS provides for the adapter selected.

When merge is supported, you can select merge actions for records that exist and records that do not exist:

If the record exists, select from:

- Update the existing record. This is the default value
- Reject the record.
- Delete the existing record.

If the record does not exist, select from:

- Include the record. This is the default value.
- Reject the record.

- Slowly Changing Dimensions (SCDs).** If the target synonym has columns with Type II SCDs, those columns will be loaded using the rules established for Type II slowly changing dimensions. If the target synonym has columns with Type I SCDs, you can choose whether to update all rows (the default) or only active rows. Although this choice appears for both types of SCDs, it is ignored for Type II SCDs.

You can re-run an SCD flow for a missed day, after subsequent days have already been processed.

When the Data Flow has an SCD target, the Load Options dialog box has the option *Overwrite Begin Date for active and End Date for inactive records*.

Check this box to select Begin Date and End Date options. Both Begin Date and End Date must be described with the same data type. The choices for overwriting the dates for Begin Date and End Date follow.

Begin Date. A Calendar Control opens, so you can select a beginning date.

End Date. The choices for End Date depend on the data type:

- DATE format (for example, YYMD).** The options are:

`The same as Begin Date for active record`

`Day prior to Begin Date for active record`

- ❑ **Date-Time (Timestamp) format (for example, HYYMDS).** Only one option is available:

`The same as Begin Date for active record`

- ❑ **Change Data Capture (CDC).** When a table is log-enabled, you can use the load option *Change Data Capture* along with the log file to update only the rows that have changed since the last load.

In order to use CDC on a Data Flow, you need a synonym for the log file, a synonym for the existing target table, a Data Flow, and, optionally, a synonym for the source table (which is not used in this type of load). When changes are made to the source table, the log file is updated. If you submit a flow that has Change Data Capture as its load option, the target table will be updated using the information in the log.

1. To create a synonym for the log, right-click the connection for the adapter with the log-enabled tables, and select *Show DBMS Objects*. On the *Create Synonym* page, select *Table Log Records* from the *Object Type* drop-down list, make any changes and selections needed for creating the synonym, and click *Add*.
2. Create a data flow using the synonym for the log file as the Source, and the existing target as the Target.
3. Right-click the Target and select *Load Options*, then select *Change Data Capture* from the Load Option drop-down list, and click *OK*.
4. In the Merge Editor, correct any missing matches (indicated in red) by dragging source columns to the Insert Expression field for the target, and click *OK*.
5. Save the flow.

If the flow is scheduled to run at set intervals, once the source table and log have been updated, the target will be updated the next time the flow runs. You can also right-click the flow and select *Submit*.

You can see the number of changes that were applied by right-clicking the flow, clicking *Logs*, and selecting *Last Log*.

- ❑ **Change Data Capture/Slowly Changing Dimensions.** Is used to load a source file that supports Change Data Capture (CDC) to an existing target that has Slowly Changing Dimensions (SCD).

When you use a source log file for an existing relational table in a flow (the Access File contains the attribute `DATA_ORIGIN=DBMSLOG`), only records that have changed are passed to the load process.

When an existing target has Slowly Changing Dimensions (the Access File contains SCD_TYPE fields), Type I records are overwritten, and Type II records are processed using either the Activation Flag or End Date/Begin Date columns to track the history of the changed rows.

- Allow Direct Load to Parquet File.** Transfers a Parquet file created in the local file system directly to the Hadoop Distributed File System (HDFS) in Parquet format, without first creating an intermediate text file. This option is available when using the Adapters for Apache Drill, Hive, and Impala.
- Create Avro Files.** Using the Adapter for Apache Hive, files can be created in the Avro self-describing binary file format. Avro is a portable (language neutral) format that can be created in the Hadoop Distributed File System (HDFS).
- Direct Load of AWS S3 Data to PostgreSQL.** If you have configured both the Adapter for AWS S3 and the Adapter for PostgreSQL, you can create a direct load flow from a delimited file in the AWS S3 repository to a PostgreSQL target.

When you configure the Adapter for AWS S3, you map the AWS S3 repository to an application on the server. Then, you can drag a file from that application directory to the Data Flow canvas as a source for the flow.

In the Load Options dialog box, select *PostgreSQL* from the *Adapter* drop-down list, select the *Bulk Load* check box, and select *On* from the *Allow Direct Bulk Load* drop-down list.

When you run the flow, the AWS S3 data will be loaded directly to the PostgreSQL table.

Adapter

Is the target adapter. Select an adapter from the drop-down list.

If no relational adapter is configured, the default Adapter will be Delimited Files.

If the Adapter for Delimited Files is configured, you can select it as the default target adapter. Navigate to the *Workspace* page. Click *Settings*, then *Settings for Reporting Server browser interface Preferences*. Scroll down to the *Target Defaults* section, and select the connection for the Adapter for Delimited Files in the *ETL-TRG-DBMS* drop-down list

If the Adapter for Excel is configured, Excel can be selected as the default adapter, as well.

Connection

Is the target connection. Select a connection from the drop-down list.

Synonym Application

Is the application name for the synonym. Accept the default application, enter an application name, or click the ellipsis (...) to select an application.

Synonym

Is the target synonym name. Accept the default synonym name, or enter a new name.

Table Name (for an SQL target)

Is the table name in the target adapter. Accept the default name or enter a name.

Bulk Load (for an SQL target)

Appears if the adapter you selected supports extended bulk load. Is checked by default. Uncheck to disable bulk load for this flow.

Save Data Files (for Bulk Load targets)

When **Bulk Load** is selected for a target in a Data Flow, the option *Save Data Files* appears.

- The default value, *No*, deletes intermediate *.ftm files from the directory for temporary files (for example, edatemp) after the corresponding portion of data is successfully committed to the target table during the Bulk Load Flow execution.
- If you select *Yes*, intermediate *.ftm files are deleted from the temporary directory (for example, edatemp) after the job completes execution.

Allow Direct Bulk Load (for specific SQL targets)


If Extended Bulk Load is used and Optimize Load is enabled, data is loaded directly to the target database without the additional extract and transformation steps, provided the following conditions are met:

- The data source is a delimited flat file.
- The file attributes including Header, Enclosure, Delimiter and Record Delimiter are supported by the target database.
- Only insert operations are used (no updates).
- No filters are used.
- No aggregation is specified.
- No transformations are used.

Key columns derived from (for an SQL target)

Choose from the following options for defining the key columns in the target.

- Key columns in source table/sorted columns in query.** The key columns will be the key columns from the source, sorted by their order in the generated query that loads the target. This is the default option.
- No key.** Creates the target with no key.
- Specify key columns.** When you click this, the first column is selected by default. You can change this by clicking the ellipsis next to this field to open the Key Editor, as shown in the following image.



Drag the columns that will comprise the key to the target frame and click **OK**.

Data File Location (for a delimited, DATREC, JSON, or Excel target)

Is the application name or location on the file system for the data file. Accept the default application, enter an application name, or click the ellipsis (...) to select an application or location on the file system.

Data File Name (for a delimited, DATREC, JSON, or Excel target)

Is the name of the target data file.

Field Delimiter (for a delimited target)

Select a delimiter from the list, or click type-in delimiter to enter a delimiter. The default delimiter is comma (,).

Header Row (for a delimited target)

When checked, adds a header row with the column names. The default is not to add a header row.

Enclosure (for a delimited target)

Specifies whether alphanumeric field values are surrounded by enclosure characters, such as single quotation marks (') or double quotation marks ("). Select *none*, if they are not enclosed, select a character from the list, or select *type-in enclosure*, to enter an enclosure character. The default value is a single quotation mark (').


3. Click *OK*.

Procedure: How to Edit the Merge Conditions for an Existing Target (Optional)


If you have an existing target and are adding source data, you can open the Merge Editor to edit the merge conditions.

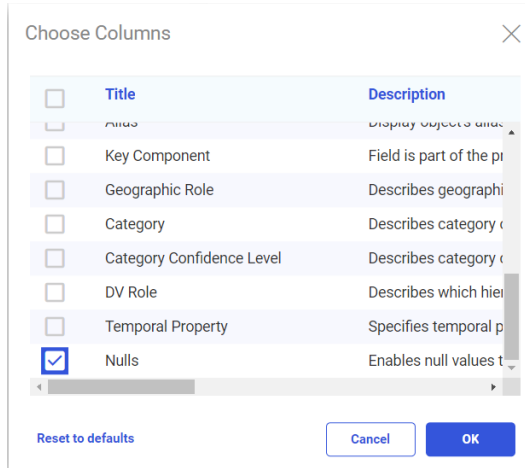
1. Right-click the target and click *Merge Editor*.

The Merge Editor opens, as shown in the following image.

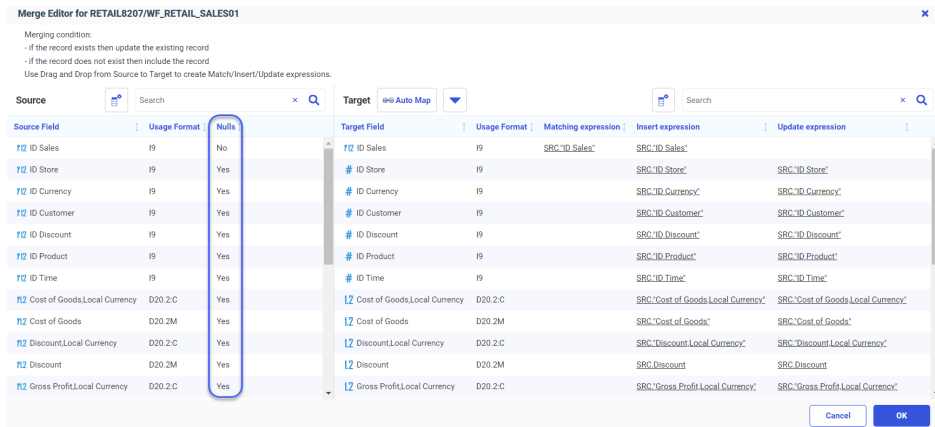


The merge strategy that was specified in the Load Options dialog box is displayed at the top. Automatic merge expressions are generated.

To see whether the columns support null values, click the *Choose Columns* icon () and select the column named *Nulls*, as shown in the following image.



After you click *OK*, you can see whether each field supports null values, as shown in the following image for the Merge Editor.



You can also choose the *Nulls* column in the *Select* and *Synonym* Editors.

2. Edit the merge expressions, if needed.

You can drag a field from the source pane to the *Matching expression*, *Insert expression*, or *Update expression* columns in the target frame to define how to match, insert, or update the row for that field.

You can right-click any underlined field to edit the expression in that column for that field, clear the expression, or clear the row. Using the *Edit Update* option in the *Update expression* column, you can, for example, add a current value to the year-to-date calculation, by adding the source field value to the existing target field value.

Click *Auto Map* to regenerate automatic merge expressions for those expressions that you cleared.

3. Click OK when you have made the needed edits.

When you run the flow, the merge will use the expressions configured in the Merge Editor.


Procedure: How to Copy and Paste Objects Between Data Flows

From the Data Flow canvas in the Reporting Server Browser Interface, you can copy one or more objects from an existing flow and paste them into another flow.


The copy and paste feature is useful if you need to reuse objects between multiple flows, for example, code (SET variables, Selects on Source, and so on). This feature is less error prone and you can save time working on the design of many flows.

Note: For a data flow, objects are copied together, except for the target object. Paste creates a new object, except when pasting an SQL object on an empty SQL object.

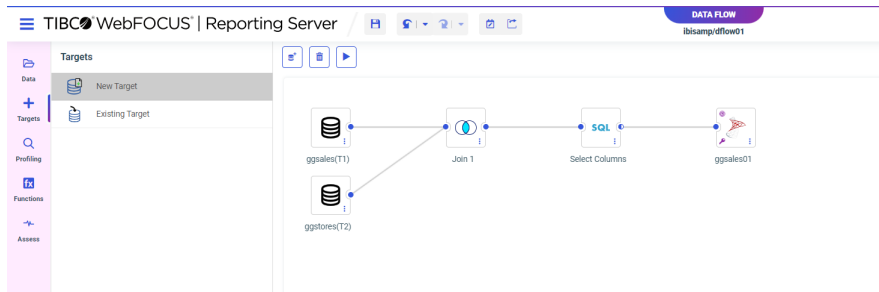
1. From the plus menu, select *Create Data Flow*, as shown in the following image.



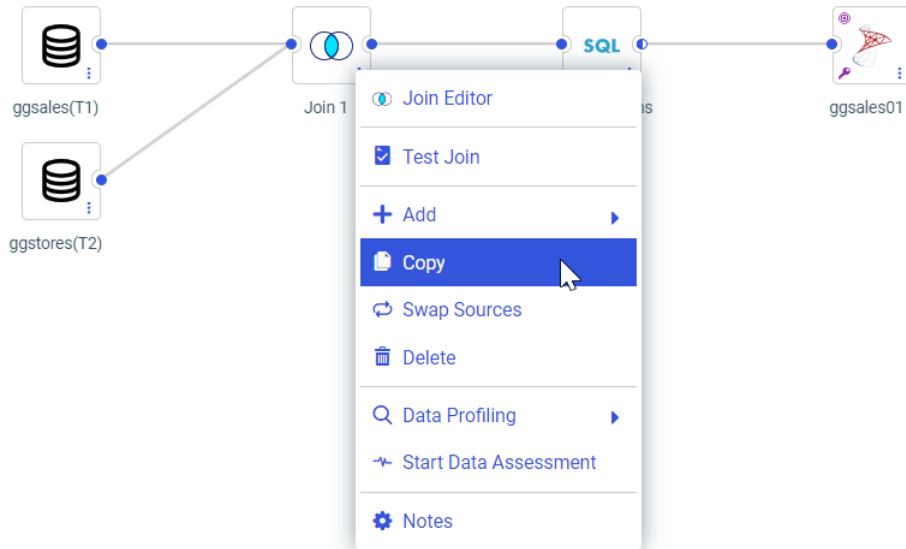
The Data Flow tab opens, as shown in the following image.




2. Add data sources and targets from the Data pane to the data flow, as shown in the following image.




3. Right-click an object in the flow and select *Copy*, as shown in the following image.



4. Create a new data flow.
5. Right-click the SQL object in the data flow and select *Paste*, as shown in the following image.



The object is pasted to the data flow, as shown in the following image.



Note: The copy and paste feature is also available in the Data Migrator Desktop Interface.

Procedure: How to Run the Flow

Prior to running the flow, make sure you have prepared the data, as described in [Preparing Data With a Flow](#) on page 376.

Click the *Run* button () to run the flow.

You can also establish a schedule for running a flow and emailing flow reports. For information, see [Running a Stored Procedure](#) on page 471.

Note: In order for the flow to run on schedule, it must be saved in a directory in the application path of the scheduler (the user id set to `sched_scan_id`). For more information, see [How to Manage Tenants](#) on page 739.

Preparing Data With a Flow

Before you load data into your target data source, you can enhance it for your needs. This massaging of the data prior to loading is called *data preparation*. You can take advantage of the data preparation options to, for example, convert numeric codes to meaningful attributes, discard erroneous data, smooth out ragged data into manageable bins, and blend descriptive data from additional data sources.

The data preparation calculations that create new fields use some of the same tools that are available in the Metadata Editor. However, when the calculations are added in the Metadata Editor, they are added to the synonym. The data remains the same, but the calculations are performed every time the field is referenced in a request against the synonym. When you create new fields in a data flow and then run the flow, the calculated values are loaded into the target data source as field values, not calculations.

Many of the data preparation examples in this topic use data from a station-based bike-share system. Bikes are unlocked from one station and returned to any other station in the system. The data we will use has been extracted from the daily ridership and membership information publically available from Citi Bike®, a public-private partnership between New York City and Lyft® Bikes.

Generating Sample Files

Citi Bike provides data monthly as zipped comma-separated values (.csv) files that contain the following data values:

- Trip Duration (seconds)
- Start Time and Date
- Stop Time and Date
- Start Station Name
- End Station Name
- Station ID
- Station Latitude/Longitude
- Bike ID
- User Type (Customer = 24-hour pass or 3-day pass user. Subscriber = Annual Member.)
- Gender (Zero=unknown; 1=male; 2=female)
- Year of Birth

To download a ridership file:

1. Go to <https://www.citibikenyc.com/system-data>.
2. Click the link that says *Download Citi Bike trip history data*.
3. Click *201907-citibike-tripdata.csv.zip* to download the file.

As part of the data preparation, this data will be augmented to have:

- Trip duration in minutes.
- Age in years.
- Alphanumeric gender values.
- Additional date components.

- ❑ Start Station Zip Code, City and County.

Depending on which zip file you download, your results may vary from the ones in this topic, which uses the data from July, 2019.

Once you have downloaded a file, you can upload it to the server using the instructions in [Uploading Files](#) on page 256.


A supplementary file was created that has station zip codes and counties (station_zip.csv). This file is attached to this PDF for your reference. After you save the attached file to your local file system, you can upload it to the server.

Enabling Sampling for a Data Flow


When a data source in a flow has a large volume of data, you can enable sampling for better response time. You can make decisions based on a sample, provided that sample is representative of the entire data set. Data Prep has a built-in capability to automatically generate a random sample (with a 99% confidence level and +/- 1% margin of error).

To enable sampling for a flow:

1. Create a flow by right-clicking an application, clicking *New*, then *Flow*, or by clicking + *Data* on the Reporting Server browser interface ribbon, and then clicking *Flow*.
2. Click *Advanced* on the ribbon, then click *Enable Sampling*, as shown in the following image.



Once you have enabled sampling, you can see if a file was sampled, and the sample size, by hovering over the file, as shown in the following image.



Data sources in a flow are represented by database-specific icons so that you can easily identify the types of sources in the flow. A generic icon is only used when a source synonym references multiple tables with disparate sources.

If a data source in a flow is not large enough to require sampling, it will not be sampled, but will be staged in the same staging target as the large sources, to eliminate joining disparate data sources and, therefore, improve join performance. By default, the staging target data source is a DATREC file, but you can configure the staging target to be the same relational data source as the load target by checking *Use ETL-TRG-DBMS for Sampling* in the Data Flow parameters of the Advanced Options dialog box that opens when you click *Options* on the data flow menu.

When sampling is enabled in a data flow and any errors occur reading the data, a warning indicator in the form of an exclamation point (!) within a red circle displays on the synonym node that has the sampling errors. Hovering over the indicator shows the messages, or a partial list, depending on the length and number of the messages.

To open a window with the complete list of messages, you can right-click the synonym and click *Show Errors*.

You can also enable sampling for all flows, so that a sample is taken automatically if needed, when the flow is opened. Set the parameter *Enable Sampling* to *On* in the Data Assist (Representative Sampling) section of the *Settings for Reporting Server browser interface Preferences* page that you open from the *Settings* menu of the *Workspace* page. For more information, see [Displaying Settings for Reporting Server Browser Interface Preferences](#) on page 601.

Using Stratified Sampling

When preparing data and viewing the distribution profile of the data, Representative Sampling can be used. This provides better responsiveness, especially for large data sets.

However, there may be times when you want to ensure that the sample includes every distinct value of one or more fields. For example, if a file has a field with the value State, and some states have very few records, a representative sample may not include all of them.

With Stratified Sampling, you can select one or more fields, and the sampling process will include records with every unique value of the selected fields, as long as the number of values does not exceed the sample size.

When sampling is enabled, this option can be selected by right-clicking a source node and clicking *Recreate Stratified Sampling*. The Stratified Sampling dialog box opens. Select the fields to use from a list of available fields, and move them to the list of selected fields using the arrow buttons. The list shows the count of unique and total values for each field. Key fields are omitted from this list, as are floating-point fields and timestamp fields with precision of less than a second, as they cannot be used.

The following image shows that the `start_station_id` field is selected for stratified sampling. This ensures that every start station will be represented in the sample, even those that are rarely used and may be omitted from a random sample.

Select Stratified Sampling Columns
✕

All columns

<input type="checkbox"/>	Name	Sampled Count	Distinct Count	Total Count
<input type="checkbox"/>	tripduration	2,742	9,857	1,047,687
<input type="checkbox"/>	bikeid	8,826	13,950	1,047,687
<input type="checkbox"/>	birth year	74	101	1,047,687
<input type="checkbox"/>	gender	3	3	1,047,687
<input type="checkbox"/>	start station na...	772	782	1,047,687
<input type="checkbox"/>	start station GI...	772	782	1,047,687
<input type="checkbox"/>	end station id	779	789	1,047,687
<input type="checkbox"/>	end station name	779	789	1,047,687
<input type="checkbox"/>	usertype	2	2	1,047,687
<input type="checkbox"/>	end station GIS ...	779	789	1,047,687

Selected columns


<input type="checkbox"/>	Name	Sampled Count	Distinct Count	Total Count
<input checked="" type="checkbox"/>	start statio...	772	782	1,047,687

>
<

Cancel Apply

When the required fields are selected, click *Apply*.

When Stratified Sampling is applied, the label of the node says *Stratified Sampling*, as shown in the following image.



To remove stratified sampling and return to representative sampling, right-click the node and click *Recreate Sampling*.

Prioritizing Sampling on Joined Tables

When Representative Sampling is enabled, you can prioritize joins in the sampling by setting the parameter *SMPL_PRIORITIZE_JOIN* to *Yes*.

When preparing data using very large data sets, the performance of many tasks can be greatly improved by working with a representative sample instead of the entire data set.

However, when data sets are joined, sampling can be problematic when not all values of a joined column are present in both samples.

For example, when joining from a (large) sales fact table to a foreign key in a customer dimension table, sampling each table separately may not include rows from the customer table that match every customer key in the sample of the fact table.

When the parameter is enabled, the first (fact) table selected is sampled first. Then, when the sample of the second (dimension) table is taken, every customer that was selected in the fact table is included in the sample of the dimension table. This ensures that every row in the fact table is matched.

To prioritize joins when sampling, do the following:

1. Navigate to the *Workspace* page of the Reporting Server browser interface.
2. Click *Settings*, then *Settings for Reporting Server browser interface Preferences*.

The *Change Settings for Reporting Server browser interface Preferences* page opens.

- Expand the *Data Assist (Representative Sampling)* section, as shown in the following image.

Data Assist (Representative Sampling)

ENABLE_SAMPLING ⓘ

 Controls Representative Sampling in Data Flow

SAMPLING_APPROOT ⓘ

 Application for representative samples. Default: foccache

SHOW_SAMPLING_INFO ⓘ

 Show Sampling Info

SMPL_ETL_TRG_DBMS ⓘ

 Use ETL-TRG-DBMS for Sampling

SMPL_PRIORITIZE_JOIN ⓘ

 Prioritize join when creating sampling

- Select *On* for *ENABLE_SAMPLING*. This enables sampling wherever it is available.
- Select *Yes* for *SMPL_PRIORITIZE_JOIN*.

This ensures that the sampling for a dimension table in a join will contain all of the key values from the sample rows generated for the fact table to which it is joined. The default value is *No*.

You can also change the following parameters.

SAMPLING_APPROOT

Is the application where the sampling files will be generated. The default is *foccache*.

SHOW_SAMPLING_INFO

The default value is *Yes*. This shows sampling information on the data flow canvas.

SMPL_TRG_DBMS

Use the adapter set as *ETL_TRG_DBMS* for sampling. The default value is *No*.

- Click *Save*.

Editing the Fields in a Flow

By default, all fields in a single-segment data source, or all fields from the top segment in a multi-segment data source are automatically added to the flow. You can turn off this option in the Advanced Options dialog box.

To edit the fields in the flow, right-click the SQL object, and click *Edit*. The Metadata and Query panes open.

You can add fields to the query or move them to a different part of the query by dragging them to the Order By, Columns, Order Across, and Filters and Variables categories in the Query panel. You can delete a field by right-clicking it and selecting Delete. In addition, you can add all of the fields from a segment to the Columns or Order By category by right-clicking the segment name, and clicking *Add to Query* and either *Column* or *Order By*.

A report of the field values opens.

You can also aggregate the data in groups by one or more field values. To perform aggregation, click the menu icon for *Order By* and select *Switch to Group by*. Drag the columns you want to group by to the *Group By* field.

Each column in the order by field is assigned a default aggregation, Sum for numeric fields, and Max for others. You can change this by selecting the column and, from the context menu, *Aggregation*, then the type of aggregation to be performed, such as Sum, Count, or Average.


Showing Profiling for a Flow

You can generate a distribution chart for each column in the flow, which shows a visual image of the field values. The profile chart displays below the field name and shows:

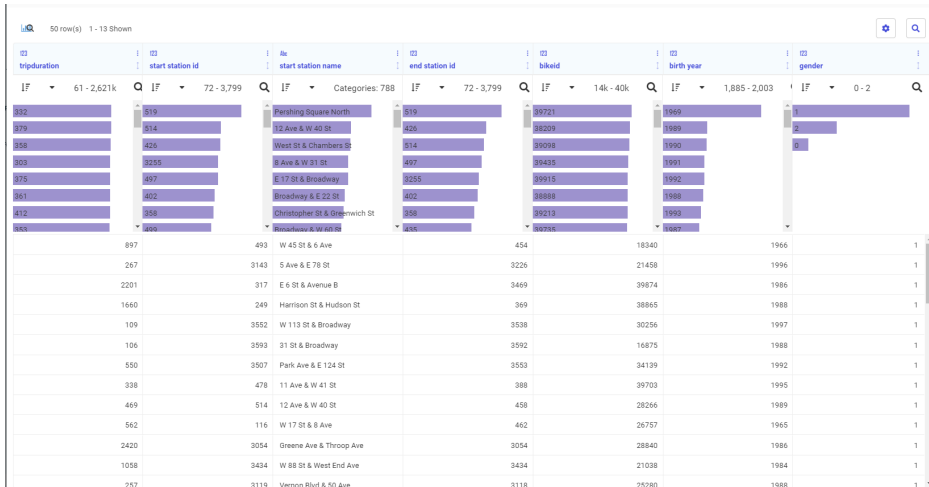
- The range of values for numeric and date fields.
- The count of unique values (categories) for character fields.

To show the profiling distribution charts:

1. Create a flow and select columns.

2. Click the *Show Profiling* button () above the sample data on the Edit Selection page or the flow page.

The profiling distribution charts display below the field names and above the field values, as shown in the following image.



Some of the interesting facts about the data shown by the distribution charts are that the number of start stations is 788 (the number of categories displayed above the distribution chart for the Start Station Name alphanumeric field), the Pershing Square North station is the most popular starting station (which makes sense as it is near a major transportation hub), and that trip durations range from 61 seconds to over 2 million seconds (a range is shown for numeric fields).



Note:

- An icon that identifies the data type of the field displays above the field name in each column. If the column was calculated, an equal sign is added before the icon.

Data Type	Icon
Alphanumeric	Abc
Integer	123
Decimal	L2

Data Type	Icon
Date-Time	
Date	
Location	


- ❑ Below the field name, there is a summary of the distribution of the data values. For numeric and date fields, it shows the range of values, and for character fields, it shows a count of categories (unique values).
- ❑ By default, the distribution charts are sorted by frequency in descending order. You can click the Sort Type button to sort them by value, and the arrow to change the sort order.

Description	Icon
Sort Type	
Sort Order	

- ❑ A search icon (magnifying glass) in each column enables you to search for values in the column. Enter the text you want to find. As you type and the search finds the value, it highlights any matching values and scrolls the chart so that the first bar with the highlighted value is at the top. To remove the search, either delete the characters or click the X to the right of the search box. To close the search box, click the X when the box is empty.

Adding or Replacing Fields

You can add new fields or replace existing fields by creating expressions.

Each field has a menu icon. 

The menu provides the following options:

- ❑ **Format.** Opens the *Edit Display Format* dialog box.

- Rename.** Enables you to change the title of the field.
- Replace with Expression.** Enables you to replace a field with values derived using a calculation.
- Add new Expression.** Enables you to create a new field with values derived using a calculation.
- Delete.** Removes the field from the flow.

Calculated fields also have:

- Properties.** Opens the Properties dialog box.
- Edit Define (Advanced).** Opens the Advanced Expression dialog box.

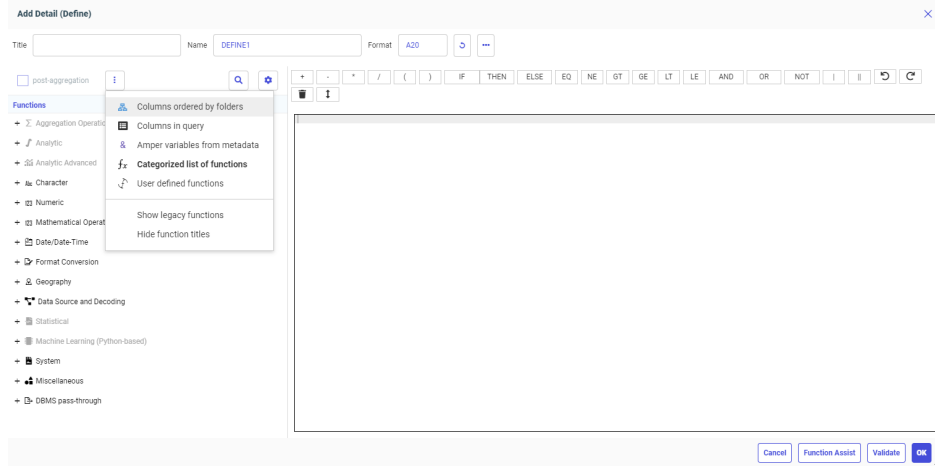
When aggregated, fields also have:

- Edit Compute (Advanced).** Opens the Advanced Expression dialog box.
- Aggregation.** Enables you to select an aggregation operator.
- Add to Filter Aggregated.** Opens a filter card for the field.

Order by fields also have:

- Sort.** Options are Ascending or Descending. The default is Descending.
- Rank.** Options are Yes (ranks the values in addition to sorting) or No. The default is No.
- Limit.** Options are No Limit, 1, 5, 10, or Custom. The default is No Limit.
- Visibility.** Options are Show or Hide. The default is Show.
- Add to Filter.** Opens a filter card for the field.

When you create an advanced expression, you use the Expression Editor. The Expression Editor provides lists of functions, columns, or variables that can be used in editing an expression. You can select which one to use from the *View list options* menu, as shown in the following image.



For more information about the Expression Editor, see [Metadata](#) on page 245.

Example: Creating a New Trip ID Field

To create the TRIP_ID field, click the menu icon in the tripduration column and click *Add New Expression*, then *Advanced Expression*.

The Add Detail (Define) Expression Calculator opens. Enter the value Trip ID for Title and Name will be filled in automatically as TRIP_ID. For the expression, enter TRIP_ID + 1. This creates a counter, where the value for each row increases by one.

Optionally, you can click *Validate* to make sure that the expression is valid. As the expression is arithmetic, the format changes from alphanumeric to integer.

Click *OK* to add this field to the flow.

Example: Creating a New Field With Trip Duration in Minutes

Trip duration is more meaningful when expressed in minutes instead of seconds.

To create the TRIP_DURATION_MINUTES field, click the menu icon in the tripduration column and click *Replace with Expression*, then *Advanced Expression*.

The Add Detail (Define) Expression Calculator opens. Perform the following steps:

1. Enter the title *Trip Duration,Minutes*. The field name becomes TRIP_DURATION_MINUTES.
2. After the text *tripduration* add a division sign and the number 60 (*tripduration / 60*) to calculate minutes.
3. Click *OK*.

The *tripduration* field is replaced by the Trip Duration Minutes field. The profiling distribution chart for this field shows that the most common trip durations are in the four to seven minute range:



Example: Creating New Fields That Calculate Time Limits and Overages

In this example, we will calculate trip limits and overages.

Trip Limit

Allowed minutes are conditional based on the user type. The allowed minutes for Subscribers is 45, and for Customers 30.

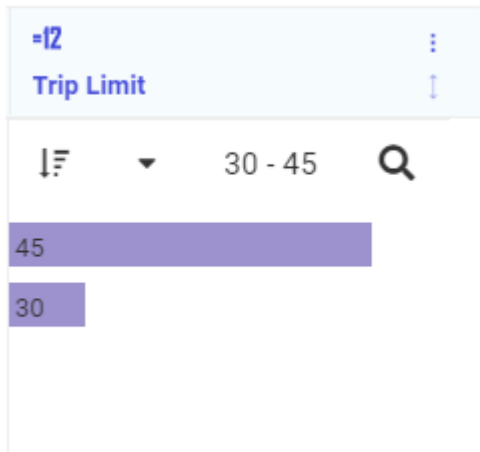
1. Click the menu icon in the *usertype* column, and click *Add new Expression*, then *Advanced Expression*.

The Add Detail (Define) Expression Calculator opens. Enter the value *Trip Limit* for Title, and Name will be filled in automatically as TRIP_LIMIT. For the expression, enter the following conditional expression:

```
IF usertype EQ 'Subscriber' THEN 45 ELSE 30
```

2. Click *OK*.

The field is added, as shown in the following image.



Base Minutes

Next, we will create a field with the number of base (included) minutes for each trip. This value is the trip duration in minutes, if that is not greater than the trip limit. If it is greater, we will use the trip limit as the value, and then calculate how many minutes of overage the trip used. We will generate the base minutes field by applying the MIN function, which returns the minimum value of its arguments.

1. Click the menu icon in the Trip Limit field, and click *Add new Expression*, then *Apply Function*.

The Apply Function dialog box opens and the Numeric folder opens.

2. Click the function *MIN - Minimum Value*.

The properties for the MIN function open, with default values for the Title, Name, Usage Format, and first expression fields.

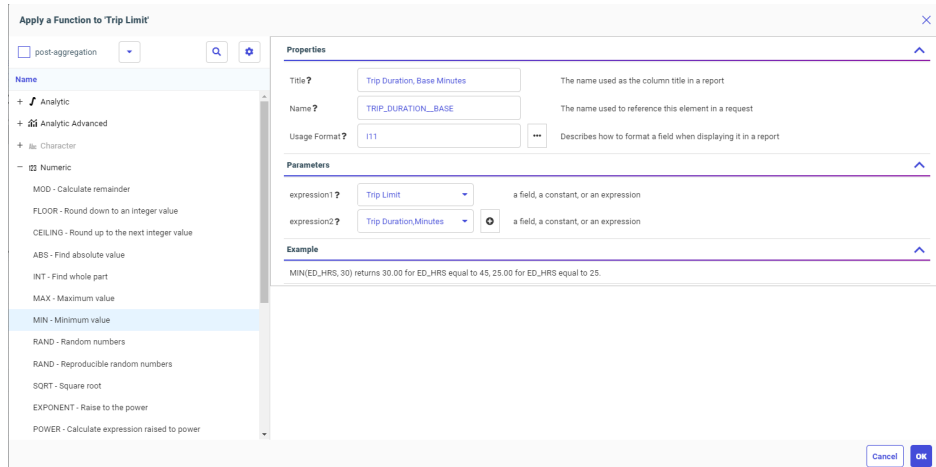
3. Change the default entries by entering the following values:

Title: Trip Duration, Base Minutes.

Name: The default name becomes TRIP_DURATION_BASE_MINUTES. Change it to TRIP_DURATION_BASE.


- ❑ **Usage Format:** I11, or click the menu button to get a properties panel where you can select Integer and type in the length.
- ❑ **Expression1:** Trip Limit (this was already there, by default).
- ❑ **Expression2:** Trip Duration,Minutes (select this field from the drop-down list).

The completed dialog box is shown in the following image.



4. Click OK.

The field is added, as shown in the following image.



Overage Minutes

Next, we will create a field with the number of minutes over the allowed amount (if any).

1. Click the menu icon in the Trip Limit field, and click *Add new Expression*, then *Apply Function*.

The Apply Function dialog box opens and the Numeric folder opens.

2. Click the function *MAX - Maximum Value*.

The properties for the MAX function open, with default values for the Title, Name, Usage Format, and first expression fields.

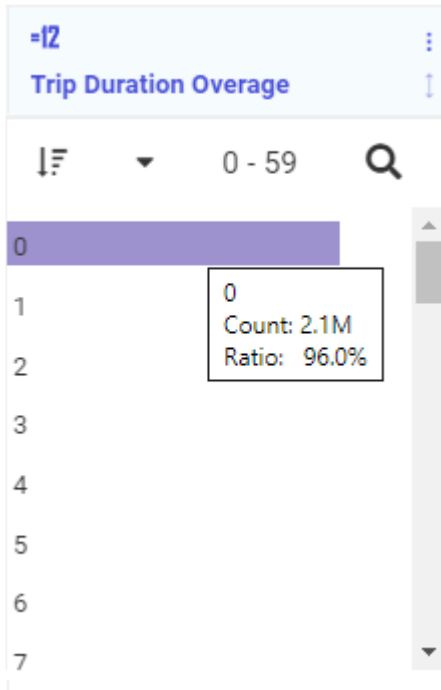
3. Change the default entries by entering the following values:

- Title:** Trip Duration,Overage. The default name becomes TRIP_DURATION_OVERAGE.
- Usage Format:** l11.
- Expression1:** 0 (zero).

- ❑ **Expression2:** TRIP_DURATION_MINUTES - TRIP_LIMIT. This calculates the overage. If it is greater than zero, this value will be returned. Otherwise, zero is returned.

4. Click *OK*.

The field is added. Since we omitted trips over 90 minutes, the overages are from 0 to 59 minutes. Hover over the zero bar and you can see that approximately 96% of the trips have no overage, as shown in the following image.



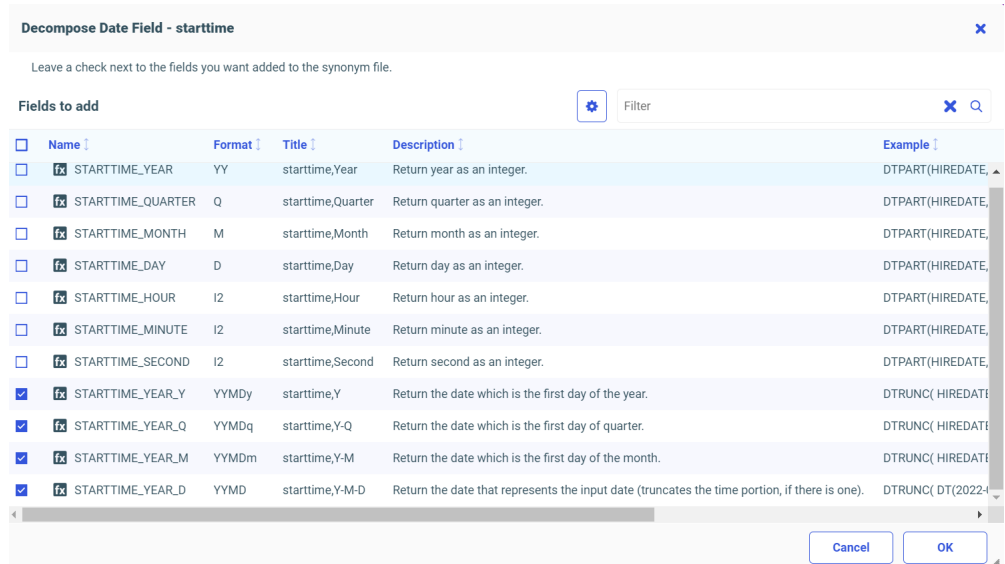
Example: Creating New Fields With Date Components

The trip data includes a timestamp, to the millisecond, that each ride started. For visualization, we want to break these fields into their component values, Year, Month, and Day of Month.

Start Year, Start Month, Start Day

1. Click the menu icon in the starttime column, click *Add new Expression*, then *Decompose Date ...*

The *Decompose Date Field* dialog box opens. It has check boxes for all the date components you can generate, as well as a description of each component and an example, as shown in the following image.



You can choose whether these descriptive items show in the Decompose Date Field dialog box by clicking the *Choose Columns* (gear) icon in the dialog box.

2. Check the following components:

- STARTTIME_YEAR
- STARTTIME_MONTH
- STARTTIME_DAY

Leave the default selections as is, as shown in the following image.

Decompose Date Field - starttime ✕

Leave a check next to the fields you want added to the synonym file.

Fields to add ⚙️ Filter ✕ 🔍

<input type="checkbox"/>	Name ↓	Format ↓	Title ↓	Description ↓
<input checked="" type="checkbox"/>	fx STARTTIME_YEAR	YY	starttime,Year	Return year as an integer.
<input type="checkbox"/>	fx STARTTIME_QUARTER	Q	starttime,Quarter	Return quarter as an integer.
<input checked="" type="checkbox"/>	fx STARTTIME_MONTH	M	starttime,Month	Return month as an integer.
<input checked="" type="checkbox"/>	fx STARTTIME_DAY	D	starttime,Day	Return day as an integer.
<input type="checkbox"/>	fx STARTTIME_HOUR	I2	starttime,Hour	Return hour as an integer.
<input type="checkbox"/>	fx STARTTIME_MINUTE	I2	starttime,Minute	Return minute as an integer.
<input type="checkbox"/>	fx STARTTIME_SECOND	I2	starttime,Second	Return second as an integer.
<input checked="" type="checkbox"/>	fx STARTTIME_YEAR_Y	YYMDy	starttime,Y	Return the date which is the
<input checked="" type="checkbox"/>	fx STARTTIME_YEAR_Q	YYMDq	starttime,Y-Q	Return the date which is the
<input checked="" type="checkbox"/>	fx STARTTIME_YEAR_M	YYMDm	starttime,Y-M	Return the date which is the
<input checked="" type="checkbox"/>	fx STARTTIME_YEAR_D	YYMD	starttime,Y-M-D	Return the date that represe

Cancel
OK

3. Click OK.

The fields are added, as shown in the following image.

#	starttime ↑ ↓	#	starttime ↑ ↓	#	starttime ↑ ↓
	Year		Month		Day
	2020		10		1
	2020		10		1
	2020		10		1
	2020		10		1
	2020		10		1
	2020		10		1
	2020		10		1
	2020		10		1

Note: In the Decompose Date Field dialog box, you can select multiple Date and Date-Time fields together for decomposition when editing a Synonym, as shown in the following image.

The screenshot shows the 'Decompose Date Field' dialog box. On the left, a list of fields is visible under the 'Time/Date Dimension' section, including 'Date', 'Date Time', 'Day', 'Day Name', 'Day of Week', 'Holiday', 'Month', 'Month Name', 'Quarter', 'Weekend', 'Year', 'Year Month', 'Year Quarter', and 'Year Week'. The main dialog area contains a table titled 'Fields to add' with the following columns: Name, Format, Title, Description, and Example. The table lists various DTPART functions such as TIME_DATE_YEAR, TIME_DATE_QUARTER, TIME_DATE_MONTH, TIME_DATE_DAY, HTIMEDATE_YEAR, HTIMEDATE_QUARTER, HTIMEDATE_MONTH, HTIMEDATE_DAY, HTIMEDATE_HOUR, HTIMEDATE_MINUTE, and HTIMEDATE_SECOND. Each row has a checkbox in the 'Name' column.

Hour of the Day

The date function DTPART returns components from a date or date-time field. Extracting the hour of the day will show how usage varies during the day.

1. Click the menu icon in the starttime column, click *Add new Expression*, then *Apply Function*.
The *Apply Function* dialog box opens, with the Date/Date-Time functions open.
2. Double-click DTPART to select it.

The Properties panel for DTPART opens.

3. Enter or select the following values.

Title: Start,Hour. The Name becomes START_HOUR.


Component: HOUR (select this from the drop-down list).

4. Click *OK*.

The field is added, as shown in the following image.




To get a better idea of usage throughout the day, click the down pointing triangle and click *Sort by Value*. The distribution chart is now sorted by hour of the day, as shown in the following image.



Example: Creating a New Field With Age

The trip data includes the birth year for each subscriber, but age would be more meaningful. The range of values for birth year starts at 1885, which is very unlikely. Also, the most frequent birth year is 1969.

Click on the bar that represents 1969. This highlights the percentage of each value in the other columns that correspond to birth year 1969. Most of them have gender 0 (not identified), as shown in the following image.



In addition, if you scroll over to the usertype column, you see that most of them are Customers (day pass buyers), not Subscribers. This indicates that 1969 is a default birth year for customers.

You can remove the selection for 1969 by clicking the red X in the column header.

To add a field for age:

1. Click the menu icon in the birth year column, and click *Add new Expression*, then *Advanced Expression*.

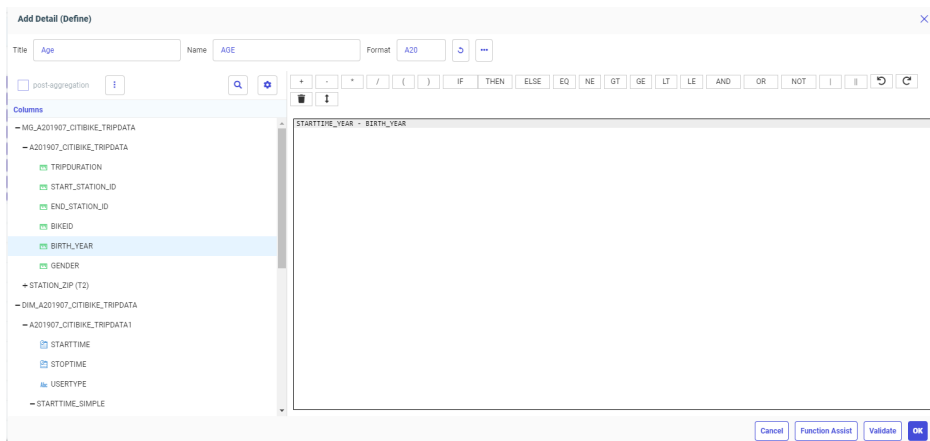
The Add Detail (Define) dialog box opens.

2. Enter the title Age. The field name becomes AGE.
3. Select *Columns ordered by folders* from the *View list options* menu.

The folder tree opens.

4. Scroll down and drag starttime,Year to the input area, enter a minus sign (-) after this, and drag birth year after the minus sign.

The following image shows this expression.



This will be the calculation for age for those riders that fall within the age parameters we want to see.

However, we want to omit riders who are Customers with the birth year 1969, or any rider born prior to 1939, so we will set their value to MISSING.

5. Edit the expression to be the following:

```
IF "birth year" LT 1939 OR ("birth year" EQ 1969 AND usertype EQ 'Customer' )  
    THEN MISSING  
    ELSE "starttime,Year" - "birth year"
```

6. Click the *Validate* button to make sure the expression is valid.

This also refreshes the format. You can also click the *Refresh* button to update the format.

7. Click *OK*.

The age range is 16 through 79, and the most common age is 30.

Example: Creating a New Field With Gender as Text

The trip data includes the gender of each subscriber as code 1 or 2 (zero for unknown). For our visualizations, we want to use the labels Male and Female, and when those are not supplied, we will use NULL or Missing.

To create the GENDER_TEXT field:


1. Click the menu icon in the gender column, and click *Add new Expression*, then *Create Decode*.

A Decode card opens.

2. Enter Male for 1, Female for 2, as shown in the following image.

Code ↓	Result
0	<input type="text"/>
1	Male
2	Female



3. Click the Edit button () and enter the following values:

Name: GENDER_TEXT.

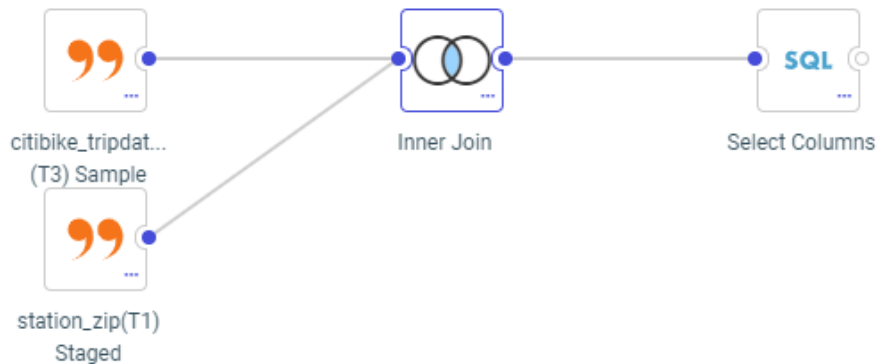
Title: Gender.

4. Click *OK*.

5. Click outside the Decode card to close it.

Join Editing

The trip data file contains station names, IDs, latitudes, and longitudes, but no zip code, city, or county information. To see the zip code and county information, the `station_zip.csv` zip file is attached to this PDF for your reference. You can upload this file to the server using the instructions in [Uploading Files](#) on page 256, and join it to the trip data file by dragging it onto the flow canvas, as shown in the following image.



The data sources in the join are represented by database-specific icons.

When you add a source to a flow as a join, by default an inner join is created based on matching fields. For tables without keys, the maximum number of automatic join fields created is three. You can edit the join to add or remove join conditions and change the type of join.

Right-click the Join object and click *Join Editor*. The Join Configuration panel shows the join properties and enables you to edit them, as shown in the following image.

Edit Join from A201907_CITIBIKE_TRIPDATA_T to STATION_ZIP_T

Configure Join ⋮

Join Type

Inner
 Left Outer
 Right Outer
 Full Outer

Join Clauses


citibike_tripdata (T1)		station_zip (T2)	
start station id	=	STATION_ID	
<input type="button" value="+"/> <input type="button" value="+ Expression"/> <input type="button" value="💡 Suggestions"/>			<input type="button" value="🗑️ Clear All"/>

The Join Editor opens in an almost full-screen window. An inner join was automatically created based on the similar field names `start_station_id` in the trip data file and `STATION_ID` in the station zip file. You can edit the join condition, add additional join conditions, and change the type of join, as described in [Creating a Join](#) on page 308.

Full outer joins are supported by WebFOCUS even for those relational data sources that do not support full outer joins.

Note: If you create a custom expression that compares a field to a variable that has no default value, you will be prompted to enter a value for the variable when you run the flow.

The Join Profiler tab shows the effects of the join. There was not a perfect match between the two files, as shown in the following image.



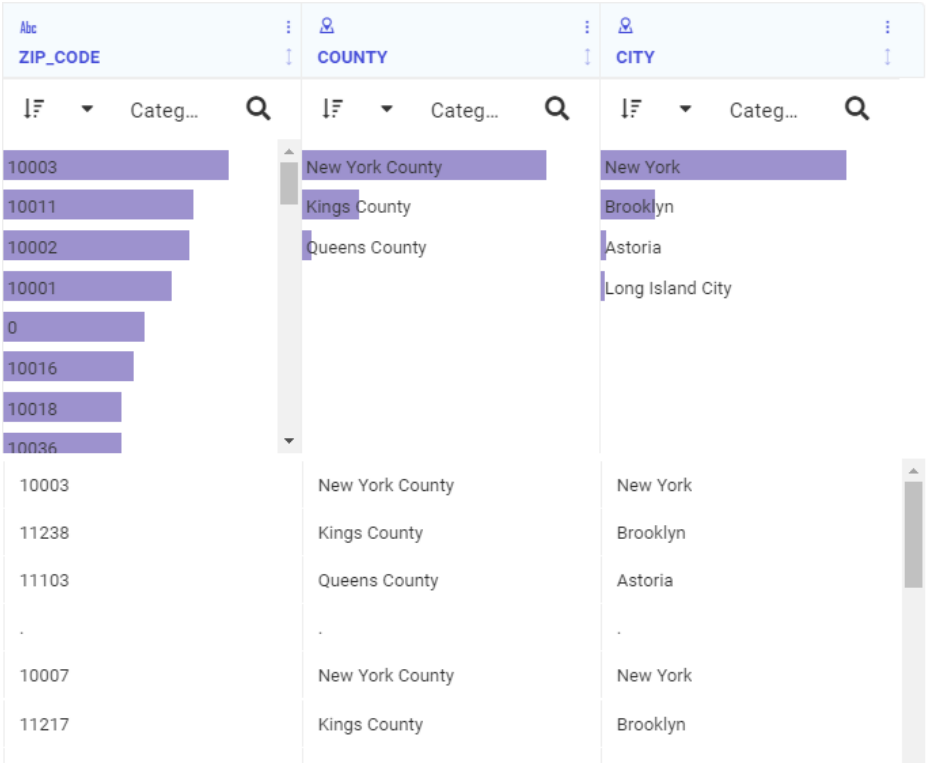
The trip data file was sampled, but the station zip file is small, so the entire file was used for the profiling. In addition, stations may have opened or closed or have been unused during the time period reflected in the data.

The gray area at the top of the first bar, and the excluded count, represent start stations that were not found in the station_zip file. For the second bar, the excluded count represents stations in the station_zip file that did not have any rides. Because of the inner join, these records will be excluded.


Click the Close button (X) at the top right of the window to return to the data flow view. Right-click the SQL object, and click *Edit*.

Multi-select the fields ZIP_CODE, COUNTY, and CITY. Right-click and select *Add to Query*, then *Column*.

Click *Show Profiling*, then scroll over to see the distribution charts and values. You will see different values because of the random sampling, as shown in the following image.



If multiple tables in a join have duplicate field names, a column name that exists in more than one table is displayed prefixed with a tag name (for example, T1) to identify which table it is from, as shown in the following image.

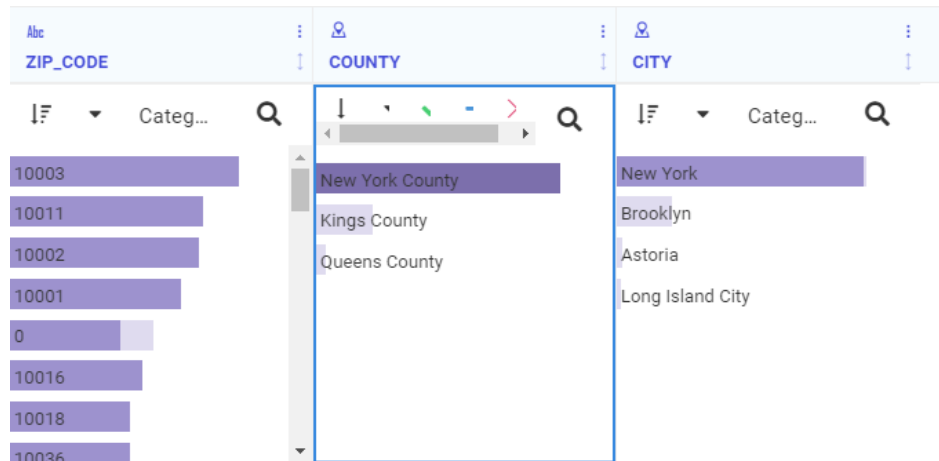


Applying a Filter

We will restrict our analysis to rides that start in Manhattan, so we can limit the data loaded to just that borough. The geographic data identifies the county, so we will just load data in New York County.

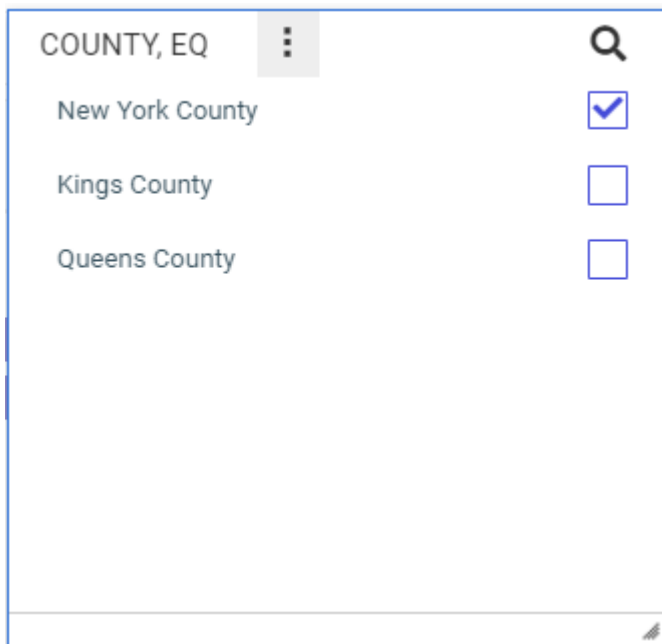
1. In the COUNTY field, click the bar for New York County.

The display changes to reflect the selection. The dark portions of the bars in each column show the proportion of rows that are selected, as shown in the following image.



2. Click the green check mark to include the selected rows.

A filter card opens, as shown in the following image.



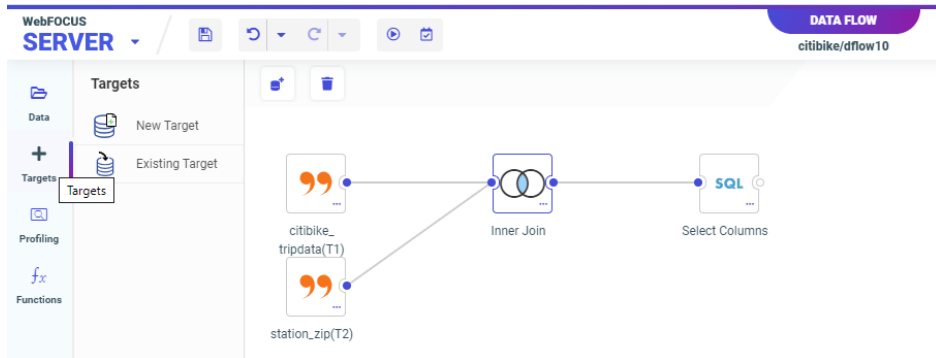
3. Click outside the filter card to close it.

Loading the Data to a Target Data Store

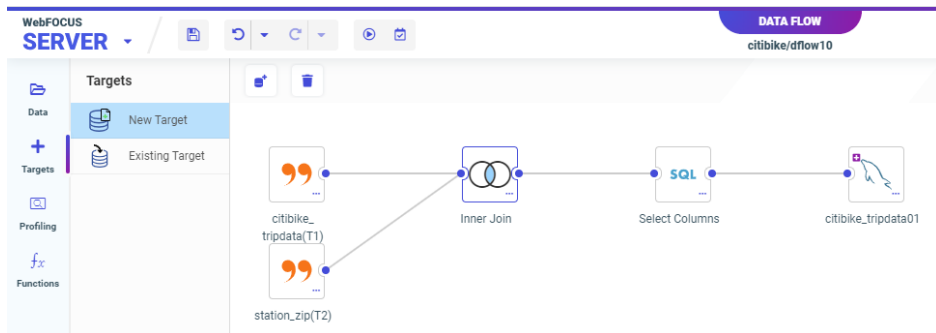
Now that the flow contains all of the fields you need, you can load it to a target data store and create a Business View.

To add a new target:

1. Click *Targets* on the sidebar of the data flow interface, as shown in the following image.



2. Click *New Target* to add a new target to the flow, as shown in the following image.



3. Right-click the target, and select *Load Options* to select or enter the target DBMS, application, synonym and table name.

For example, the following load options load the data into an Oracle table, to an application named citibike, with synonym and table name bike_share_nyc.

The screenshot shows a 'Load Options' dialog box with the following configuration:

- Load Option:** New/Replace
- Adapter:** Oracle
- Connection:** CON1
- Synonym Application:** citibike
- Synonym:** citibike_tripdata01
- Table Name:** citibike_tripdata01
- Bulk Load:** (unchecked)

Buttons at the bottom right include 'Cancel' and 'OK'.

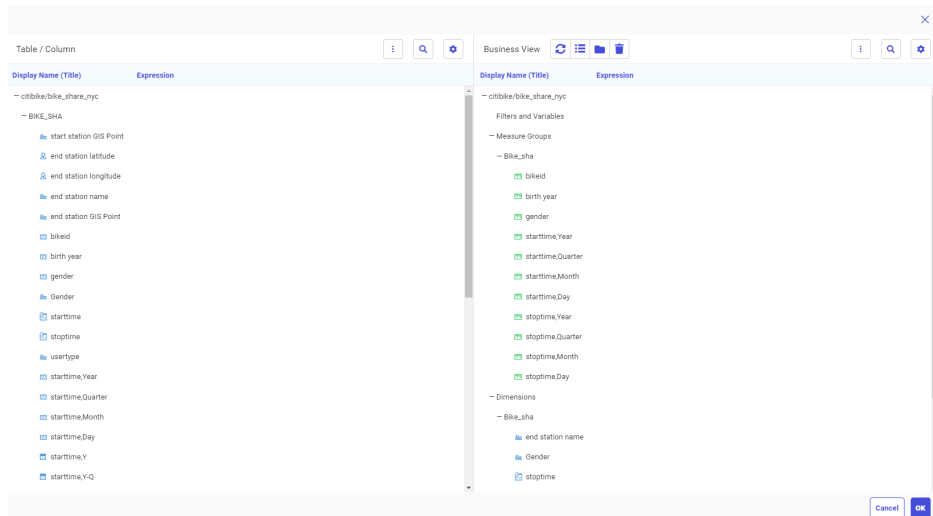
4. Click *OK*.

For more information on load options, see [How to Edit the Load Options for the Target \(Optional\)](#) on page 364.

Creating a Business View

You can create a Business View for the target to categorize the columns into measures and dimensions. This view of the data is used by WebFOCUS Designer for creating reports and charts.

Right-click the target and select *Edit Business View*. A default Business View is created, as shown in the following image.




By default, numeric fields are added as measures and alphanumeric fields are added as dimensions. You can edit the Business View using the options available. You can reclassify fields to change their DV roles, remove selected fields from the Business View, rename folders, and add new folders.

For this example, make the following edits:

- Make Trip ID a dimension field.** Right-click *Trip ID* and select *Cut*. Then right-click *Dimensions* and select *Paste*.
- Rename the start folder.** Right-click the folder named *start* and select *Properties*. Enter the name *Start Station* and click *OK*.
- Rename the starttime, Simple folder.** Right-click the folder *starttime, Simple* and select *Rename*. Enter the name *Start Date* and click *OK*.
- Move Hour and Week fields.** Multi-select *Start,Hour* and *Start,Week*, and select *Cut*. Scroll down to *Start Date*. Right-click and select *Paste*.
- Move *start station id* to the dimension *Start*.
- Move *end station id* to the dimension *End*.
- Remove the fields bikeid and trip limit from the Business View.** Multi-select those fields, right-click, and select *Hide*.

- ❑ **Create a new dimension.** Right-click *Dimensions* and click *Manage Folders*, then *New Folder*. A folder is added named *Dimension(1)*. Right-click the folder and select *Rename*. Enter the name *Rider Information* and click *OK*.
- ❑ Move the fields *usertype*, *birth year*, and *age* to *Rider Information*. At this point, the only fields left in *Measures* should be *Trip Duration,Minutes*, *Trip Duration,Base Minutes*, and *Trip Duration,Overage*.
- ❑ **Add a hierarchy to Start Station.** Right-click *Start Station* and select *Manage Folders*, then *New Folder*. A new folder is added called *Hierarchy (1)*. Right-click the folder and select *Rename*. Enter *Location* and click *OK*.
- ❑ Move the fields from *geography* to *Location*.


Loading the Data


On the toolbar, click the *Run* button ().

When the data is loaded, a report similar to the following displays.

```
(ICM18122) Request - citibike/dflow02.fex (Owner: USER1) submitted.
(ICM18741) citibike/citibike_tripdata01 type Oracle New target
(FOC2662) BULK LOAD PROCESS STARTED AT 12.05.04
(FOC2661) TARGET FILE CITIBIKE/CITIBIKE_TRIPDATA01
(ICM18745) Commit forced at: 1000000 for 1000000 row(s)
(ICM18745) Commit forced at: 2000000 for 1000000 row(s)
(ICM18745) Commit forced at: 2077758 for 77758 row(s)
(FOC2663) BULK LOAD PROCESS ENDED AT 12.07.29, ELAPSED TIME = 00:02:25.219
(FOC2661) TARGET FILE CITIBIKE/CITIBIKE_TRIPDATA01
CITIBIKE/citibike_tripdata01 HELD AS SQLORA TABLE
1
0 NUMBER OF RECORDS IN TABLE= 2077758 LINES=2077758
0
(ICM18744) Ending Load
(ICM18040) Return Code = 0
(ICM18076) Request: citibike/dflow02.fex - finished processing
(ICM18007) CPU Time : 106125
```

Editing Data Flow Options

The Data Flow workspace has an advanced options menu (), as shown in the following image.



The following options are available.

Submit. Submits the flow to run in the background, with output returned to a log file.

Run with Limited Data. Runs the flow with a limited number of rows (by default, 50), for diagnostic and testing purposes.

Logs. Gives you the following options:

- Last Log.** Opens a pane showing the most recent log.
- Last Output.** Opens a pane with the most recent flow output.
- Log and Statistics.** Opens a pane with a dialog box for running a report of logs or statistics, as shown in the following image.

The screenshot shows a dialog box with the following fields:

- Report Type?: Statistics
- Report Range?: All
- Completion Status?: All
- Run by?: All
- Application or Flow Name?: citibike/dflow10.fex

Buttons: Cancel, View Report

The choices are:

Report Type

Select *Statistics* or *Log*. The default is *Statistics*.

Statistics reports display statistics about the flow, such as user ID, application, log file, start and end date and times, run times, targets, return codes, transactions, accepts, messages, and other relevant statistics.

Log reports display data flow output.

Report Range

Select *Last*, *All*, or *Specify Time*. The default is *All*. *Specify Time* opens date and time pickers for the start and end of the range.

Completion Status

Select *All*, *Completed*, *Completed Successfully*, or *Completed With Errors*. The default is *All*.

Run by

Select *All*, *Scheduler*, or *Submit*. The default is *All*.

Application or Flow Name

Enter the application and flow name, accept the application and flow name already in the entry field, or click the ellipsis (...) to navigate to the application and flow for the report.

When you have made your selections, click *View Report*.

Options. Opens the Advanced Options dialog box. For information, see [Advanced Options Dialog Box](#) on page 415.

Enable Sampling. Creates representative samples of large sources for profiling to improve response time.

Diagnostics. Provides the following option.

View Source. Displays the flow request syntax.

Clear Undo History. Clears the undo history to prevent previous actions from being undone.

Enable Traces. Turns on tracing. You can view traces by going to the Workspace page, expanding *Logs and Traces*, and double-clicking *Traces*. A list of traces opens. Double-click a trace to view it.

Session Log. Opens the Session Log, which contains a list of the commands that have been executed in the session.

Clear Session Log. Clears the session log.

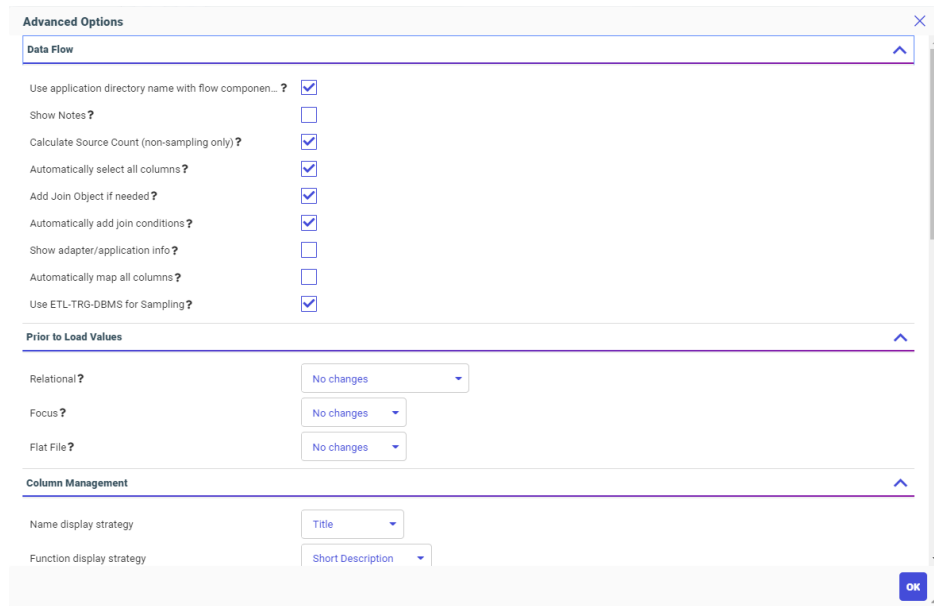
Login Info. Displays information about the connected user and resources.

Preferences. Opens the Reporting Server browser interface Preferences dialog box. For information, see [Displaying Settings for Reporting Server Browser Interface Preferences](#) on page 601.

Language. Enables you to select one of the server-supported languages.

Reference: Advanced Options Dialog Box

The *Advanced Options* dialog box is shown in the following image.



You can edit the following options.

Data Flow Options

Defines the following options for Data Flows.

Use application directory name with flow components

When selected, prepends the application directory name to each flow component, generating a two-part name in the form app/component. By default, this option is selected.

Show Notes

You can add notes in the Properties dialog box that opens when you right-click a flow component and click *Properties*. When selected, any note you add displays as a sticky note with its component on the Data Flow workspace. You can close a sticky note by clicking the Close button (X) in the upper right corner of the note. By default, this option is not selected.

Calculate Source Count (non-sampling only)

When selected, calculates the count of source rows when sampling is not enabled. The count displays when the mouse hovers over a source. By default, this option is selected.

Add Join Object if needed

When selected, adds a join object when adding a source to the workspace without specifying join or union. By default, this option is selected.

Automatically add join conditions

When selected, automatically creates join conditions when a join is added, based on similar names. By default, this option is selected.

Show adapter/application info

When selected, displays the adapter and application for sources and targets on the Data Flow workspace. By default, this option is not selected.

Automatically map all columns

When selected, automatically maps all columns when a target is added. By default, this option is not selected.

Use ETL-TRG-DBMS for Sampling

When selected, stages sample data using the same adapter and connection as the target. By default, this option is selected.

Prior to Load Values

Defines the following options for existing targets prior to loading them.

Relational

Choices are:

- No changes. This is the default value.
- Delete all rows from table.
- Truncate table.

Focus

Choices are:

- No changes. This is the default value.
- Drop Table.

Flat File

Choices are:

- No changes. This is the default value.
- Delete File.

Column Management

Defines the following options for columns.

Name Display Strategy

Defines the Master File attribute to display as column title. Choices are:

- Name.
- Title. This is the default value.
- Description.
- Alias.

Function display strategy

Defines how to display function information in the Expression Editor and Function Assist. Choices are:

- Syntax.
- Short Description. This is the default value.

Language Generation Options

Defines the following options for generating the flow request.

Use segment to qualify field reference

Defines whether to use qualified field names. Choices are:

- For duplicate fields only. This is the default value.
- Always.

Use application name to qualify synonym reference

Defines whether to use two-part synonym names (app/synonym). Choices are:

- Yes. This is the default value.
- No.

Use homeapps/user instead of myhome

Defines whether to use two-part names for the home directory of a user. Choices are:

- Yes. This is the default value.
- No.

Limits

Defines the following options for Data Flow limits.

Maximum number of rows for test

Enter a number of rows to use for the Test option for a flow. The default value is 50.

Maximum number of columns for test

Enter a number of columns to use for the Test option for a flow. The default value is 999999.

Maximum number of rows for Feature Importances test

Enter a number of rows to use for the Feature Importances data profiling chart. The default value is 2000.

Maximum number of rows for Box plot test

Enter a number of rows to use for the Box plot data profiling chart. The default value is 100000.

Limit number of database reads for a filter

Enter a number of database reads when applying a filter. The default is blank (no limit).

Maximum number of identical error messages

Enter a limit for the number of identical error messages that can be displayed. The default value is 10.

Undo/Redo Limit

Enter a number for the maximum number of actions that can be undone or redone. The default value is 50.

Number of rows in bins

Enter a maximum for the number of rows assigned to a bin. The default value is 32.

Bins for date unit

Select the data unit that should be used for binning on date fields. Choices are:

- Auto. The unit is selected automatically based on the input date format. This is the default value.
- Year.
- Quarter.
- Month.
- Week.
- Day.

Omit missing (null) values in Data Profiling charts

Select Yes to omit missing values when generating data profiling charts. The default value is No.

Data Tree Display Options

Defines the following option for listing applications on the data pane:

Use Monotree

Select *No* to display the application directories as a hierarchical tree in which you click the plus sign (+) to expand a directory and the minus sign (-) to collapse it.

The default value, *Yes*, displays the application directories as closed folders.

When you click a folder, that folder expands and the other folders are not displayed. Click again to return to the original list of folders.

Profiling Sort Options

Defines the following option for sorting profiling columns.

List included with temporary filter values first

The choices are *Yes* and *No*. The default value is *No*.

Filter

Defines the following option for filters.

Advanced Filter

Defines whether advanced filters are supported. The choices are *Yes* and *No*. The default value is *Yes*.

Advanced Debug Options


Defines the following option.

Sleep Time

Enter a number to pause processing.

Reference: Data Flow View Options

The Data Flow canvas has a View Options menu, as shown in the following image.



The following View Options are available.

Arrangement

Enables you to choose how the components of a flow are arranged on the canvas. Choices are:

- Vertically spread out.** Spreads out the components to a reasonable distance apart vertically.
- Vertically Compressed.** Compresses the vertical distance between components.
- Vertically Compressed with long H-space.** Compresses the vertical distance between components and lengthens the connectors.
- Auto Arrange.** Automatically arranges the distance between components on the canvas.

Caption

Enables you to choose between displaying the Name or Tag of the synonyms on the canvas. The default option is Name.

Icons

Enables you to choose between Large or Small icons for the components on the canvas. The default value is Large.

Flow Properties

Opens the Flow Properties dialog box. The properties displayed depend on the type of Data Flow that is open in the canvas. You can change options, such as restart options, and click *Apply* to implement them.

More options...

Opens the Advanced Options dialog box, described in [Advanced Options Dialog Box](#) on page 415.

Assigning Slowly Changing Dimension (SCD) Columns

A Slowly Changing Dimension (SCD) is a dimension in which the data does not change frequently or on a regular schedule. Slowly Changing Dimensions are common in data warehouses.

When a Synonym has SCD columns defined, the changes to the records in the database can be handled following Type I and Type II (Type 1 and Type 2) implementation techniques. The SCD implementation techniques are applied when such a Synonym is being used as a target in a Data Flow, and the Load Option on this target is set to *Slowly Changing Dimensions*.

For example, customer addresses change only when a customer moves. A customer who moves may change spending habits at the new address, so tracking customer data for that customer may depend on the address. Therefore, it could be useful to keep historical records of those changes.

Other types of changes may not require keeping a record of the change.

WebFOCUS supports two methodologies for handling Slowly Changing Dimensions.

- ❑ **Type I Processing.** SCD Type I processing does not store historical data in the dimension table. This method overwrites the old data in the dimension table with the new data. One common use for this is to correct data errors in the dimension, for example, to correct a misspelling. This type of processing is easy to maintain, and it does not add additional rows to the data.
- ❑ **Type II Processing.** SCD Type II processing stores the entire history of the data in the dimension table. With this type of processing, unlimited history can be tracked. With Type II processing, WebFOCUS supports storing the data in two different ways, which can be used separately or in combination with each other. They are:
 - ❑ **Flagging.** With flagging, an activation flag column in the data identifies whether a row is active (current) or inactive (historical). Every time a change is made, a new row is created, with the flag value identifying it as active (for example, the value 1), and the record that was just replaced has its flag column changed to the inactive value (for example, 0).
 - ❑ **Effective Date.** With effective date processing, the period of the change is tracked using `start_date` and `end_date` columns in the dimension table. The current row has NULL in the `end_date` column.

You can assign the required SCD column values in the Reporting Server browser interface.

Reference: Summary of SCD Columns

The following columns define SCD processing.

Surrogate Key (required)

Is an Integer column that uniquely identifies the row in the dimension. The Surrogate Key column must be assigned.

Logical Key (required)

Is the Source database key, which can consist of multiple columns. The Logical Key column must be assigned.

Type I

Defines columns whose values are overwritten with new values, no history preserved. Either Type I and/or Type II columns must be assigned.

Type II

Defines columns whose older values are preserved with inactive status, and whose new values are added with active status, in order to preserve the history. Either Type I and/or Type II columns must be assigned.

Begin Date

Is the Date or Date-Time column used to set the time when the row becomes the active row. Data Migrator sets a default value.

The format of this column should be identical with the format of the End Date column.

The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.

End Date

Is the Date or Date-Time column used to set the time when the row becomes an inactive row. The user can set a value, otherwise, Data Migrator sets a default value.

The format of this column should be identical with the format of the Begin Date column.

The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.

Activation Flag

Is an Integer or Alphanumeric column that indicates if the row is current. The user can set a value, otherwise, Data Migrator sets a default value.

The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.


Change Flag

Is a Date or Date-Time column that is used to set the time when the value of a Type I column changes. The user can set a value, otherwise, Data Migrator sets a default value.

***Procedure:* How to Assign SCD Columns**

1. Open the synonym that has SCD columns in the Synonym Editor, by double-clicking the synonym, or by right-clicking and selecting *Open*.

2. Right-click the segment that has the SCD columns, and click *Assign SCD Columns*, as shown in the following image.



The Assign SCD Columns pane opens.

3. Drag the appropriate fields from the segment fields list to the SCD columns in the *Assign to* pane, as shown in the following image.

Assign SCD Columns for DIMREP_SCD_M1

Select the items you wish to assign to a slowly changing dimension type and drag and drop into the desired area.

Fields

DIMREP_SCD_M1

Filter List

- Surrogate Key
- Employee ID
- Branch
- Region
- First Name
- Last Name
- Start Date
- End Date
- Active

Assign to

- Surrogate Key Field**
 - Surrogate Key
- Logical Key Field**
 - Employee ID
- Begin Date**
 - Start Date
- End Date**
 - End Date
- Activation Flag**
 - Active
- Type I**
 - Branch
- Type II**
 - First Name
 - Last Name
- Change Flag**
 - Drag a field to here

Cancel Validate Apply

You can click the Information icon in each row, if you require help with the specific column, as shown in the following image.

Assign SCD Columns for DIMREP_SCD_M1 ✕

Select the items you wish to assign to a slowly changing dimension type and drag and drop into the desired area.

Fields

DIMREP_SCD_M1

Filter List ✕ ▾

- Surrogate Key
- Employee ID
- Branch
- Region
- First Name
- Last Name
- Start Date
- End Date
- Active

Assign to

- Surrogate Key Field

 - Surrogate Key ⓘ ✕
- Logical Key Field

 - Employee ID ⓘ ✕
- Begin Date

 - Start Date ⓘ ✕
- End Date

 - End Date ⓘ ✕
- Activation Flag

 - Active ⓘ ✕
- Type I

 - Branch ⓘ ✕
- Type II

 - First Name ⓘ ✕
 - Last Name ⓘ ✕
- Change Flag

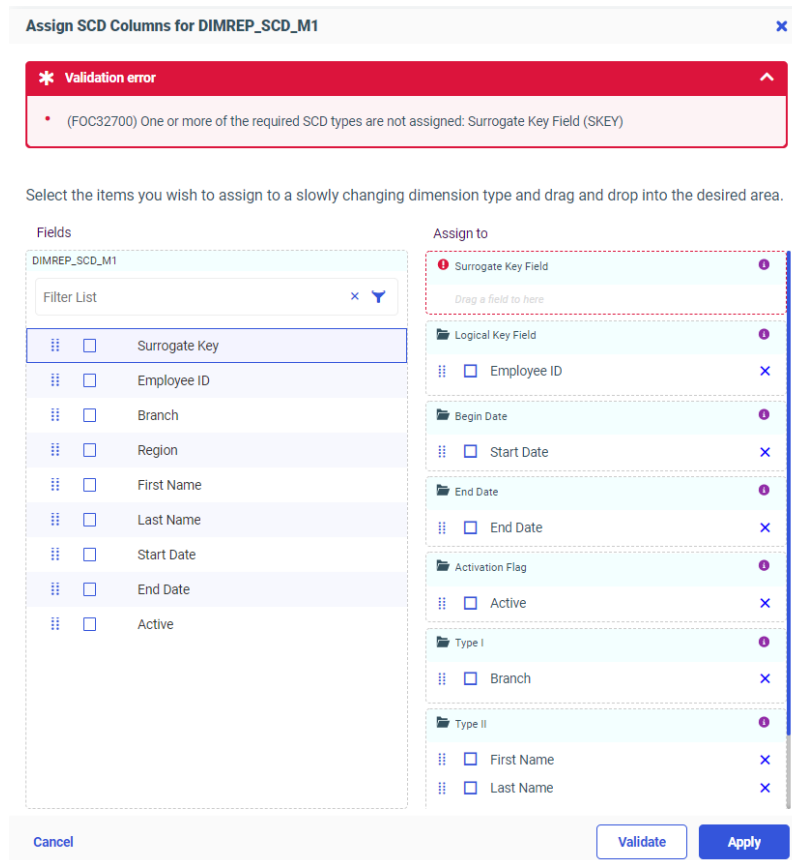
Drag a field to here

Cancel
Validate
Apply

4. When you have finished, click *Validate* to check that the columns you configured are valid.

Generating Data Flows Using a Template Flow

If there are invalid selections, messages are displayed at the top of the pane in red, as shown in the following image.



5. If necessary, correct any issues.
6. Click *Apply*.

Generating Data Flows Using a Template Flow

You can use a template data flow that has one source and one new or existing target to generate multiple new data flows. A process flow will be created that contains all of the generated data flows. This feature enables you to easily move any number of single sources to new targets.

When you use a template flow to generate new flows, on the Generate Flows screen you can change the default application directory locations, specify source and target prefixes and suffixes, change the Process Flow name, select parallel execution, and select the objects for which to apply the prefixes and suffixes. Once you click *Generate and Save* or *Generate and Submit*, these selections are saved and appear the next time you start the Generate Flows process using the same template flow.

Generate Flows can also be used with:

- A template Data Flow that has one source and one target, where the target synonym has Slowly Changing Dimensions (SCDs), and in which the Load Option is *Slowly Changing Dimensions*.
- A template flow that uses a Source Log synonym (Change Data Capture source) and one target that has Load Option *Change Data Capture*.

Procedure: How to Generate Data Flows

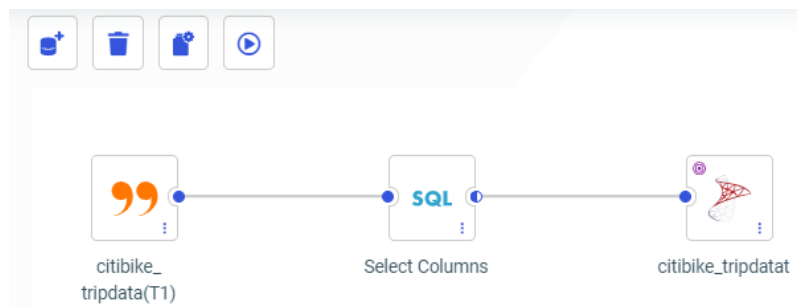
The template data flow and the generated data flows can each have one source and one new or existing target.

WHERE filters are not supported in the SQL statement in the template Data Flow.


1. Create and save a data flow that has one source and one target to use as the template.

The template flow is a normal data flow that can have any name and be saved in any application directory.

For example, the following image shows a flow named `template1` that has a delimited source and a new SQL Server target.




2. Right-click the template flow and point to *Generate Flows*, as shown in the following image.



Note: The *Generate Flows* option appears as long as the flow conforms to the requirements for a template flow. That is, it can have one source and one target, and no WHERE filter.

When a Data Flow that conforms to the requirements for a template flow is open in the flow canvas, there is a *Generate Flows* button at the top of the window, as shown in the following image.



Click either:

- Generate & Save* to save the generated flows but not submit them automatically. Once they are saved, you can run or submit them manually. This option gives you the ability to edit generated flows prior to submitting them.
- Generate & Submit* to submit the generated flows. The flows are not saved, they are generated in the temporary foccache application and deleted when the session ends. Only the target synonyms and tables are saved.

The *Generate Flows* window opens, as shown in the following image.

Generate Flows based on templateflow/template1

Configure flow | Template information | Schedule information

Source

Application Folder

Prefix Suffix

No Synonyms Found
Select different source application.

New target

Application Folder

Prefix Suffix

Apply Prefix and Suffix To

Flow


Save Location

Process Flow Name

Parallel execution

3. Select the source, target, and, if you are saving the flows, the flow application folders.

If the source application folder does not have the sources you want to use, click the application name and navigate to the application you need, as shown in the following image.



Do the same for the New or Existing Target application and the Flow application (if you are saving the flows). The source and target applications must be different, if the target names match the source names. The Target application will contain the target synonyms. The Flow application will contain the generated data flows and process flow. The Target and Flow application folders can be the same or different.

4. Optionally, type over the default process flow name to generate the process flow name you need.
5. Optionally, select the *Parallel execution* check box to create a group with the data flows inside.

6. Select the sources for the new flows, as shown in the following image.

Generate Flows based on templateflow/template1

Source

Application Folder:

Prefix: Suffix:

New target

Application Folder:

Prefix: Suffix:

Apply Prefix and Suffix To:

Flow

Save Location:

Process Flow Name:

Parallel execution

AVAILABLE SOURCES

<input type="checkbox"/>	Source Name	Target Name	Target Table Name
<input type="checkbox"/>	database_of_company	database_of_company	database_of_company
<input checked="" type="checkbox"/>	dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/>	dmhr	dmhr	dmhr
<input type="checkbox"/>	dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/>	dminv	dminv	dminv
<input type="checkbox"/>	dminv_log	dminv_log	dminv_log
<input type="checkbox"/>	dminva	dminva	dminva
<input type="checkbox"/>	dminvd	dminvd	dminvd
<input type="checkbox"/>	dminvf	dminvf	dminvf

Cancel Generate & Save

The sources do not have to be the same type as the source in the template flow. Any source type can be selected.

7. New Target Instructions

Optionally, use one of the following techniques to edit generated new target synonym and table names.

- You can type over the default target names to create new names.
- You can use the Source *Prefix* text box to specify characters that will be removed from the start of the source synonym names when creating the target names, and the Source *Suffix* text box to specify characters that will be removed from the end of the source synonym names when creating the target names.
- You can use the Target *Prefix* text box to specify characters to be added to the start of the target names, and the Target *Suffix* text box to specify characters to be appended to the end of the target names.

Generating Data Flows Using a Template Flow

By default, the Prefix and Suffix transformations will be applied to both target synonym and table names. However, you can use the *Apply Prefix and Suffix To* drop-down list to select whether to apply them to:

- Target Synonym and Table names (the default).
- Target Synonym names only.
- Target Table names only.

When you enter one or more prefix and suffix values and move the focus away from the entered values, the target names are automatically updated on the synonym selection panel at the bottom of the screen.

On the following Generate Flows screen, if the prefix *dm* is found in the source synonym name, it is to be removed from the target synonym and table name. The suffix *_tgt* is to be added to all target synonym and table names. The target names appear with these adjustments, as shown in the following image.

Generate Flows based on foccacha/genflow/template01

Configure flow | Template information | Schedule information

Source

Application Folder: ibisamp

Prefix: dm | Suffix:

New target

Application Folder: ibisamp

Prefix: | Suffix: _tgt

Apply Prefix and Suffix To: Target Name and Target Table Name

Flow

Save Location: generatedflow

Process Flow Name: template01_g

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
database_of_company	database_of_company_tgt	companydatabase/_search_tgt
dmcomp	comp_tgt	comp_tgt
dmhr	hr_tgt	hr_tgt
dmhrinp	hrinp_tgt	hrinp_tgt
dminv	inv_tgt	inv_tgt
dminv_log	inv_log_tgt	marina.dbo.dminv_tgt
dminva	inva_tgt	inv_tgt
dminvd	invd_tgt	invd_tgt
dminvf	invf_tgt	inv_tgt

Cancel | Generate & Save

For example, the source synonym name *dminv* will generate a target synonym named *inv_tgt*.

Optionally, you can type over the target table names to generate the table names you need.

8. Existing Target Instructions

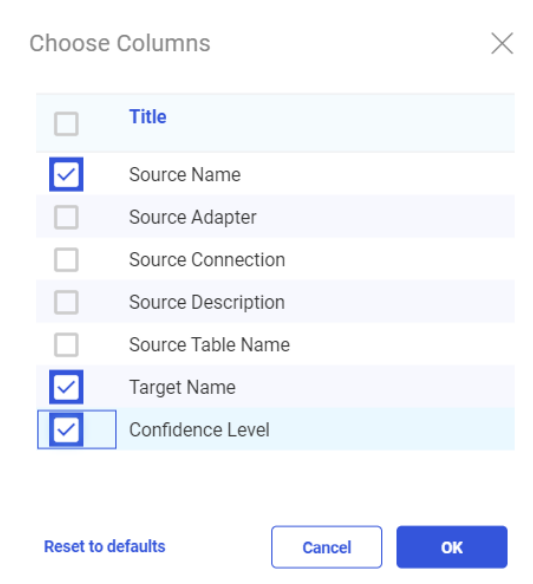
When you use Generate Flows with an existing target, the Generate Flows screen opens initially showing the default suggested target names. These suggested targets are identified by a folder icon with a white fill color, as shown in the following image.

The screenshot shows the 'Generate Flows based on foccache/genflow/template_existing' dialog box. The 'Configure flow' tab is selected. The 'Source' section has 'Application Folder' set to 'ibisamp' and 'Prefix' set to 'dm'. The 'Existing target' section has 'Application Folder' set to 'ibisamp' and 'Prefix' set to '_tgt'. The 'Flow' section has 'Save Location' set to 'generatedflow' and 'Process Flow Name' set to 'template_existing_g'. Below these sections is a table titled 'AVAILABLE SOURCES' with columns for 'Source Name' and 'Target Name'. The table lists various source names and their corresponding target names, each with a folder icon. The row for 'dminv' and 'inv_tgt' is highlighted. At the bottom right, there are 'Cancel' and 'Generate & Save' buttons.

Source Name	Target Name
dmhr	hr_tgt
dmhrinp	hrinp_tgt
dminv	inv_tgt
dminv_log	inv_log_tgt
dminva	inva_tgt
dminvd	invd_tgt
dminvf	invf_tgt
dmord	ord_tgt
dmordb	ordb_tgt

The row for each existing target has a confidence level associated with it. The confidence level is expressed as a percentage and identifies the compatibility between the source and target based on matching column names and data types.

By default, the Confidence Level column does not display on the Generate Flows screen. To show it, click the Choose Columns icon, select *Confidence Level*, and click *OK*, as shown in the following image.

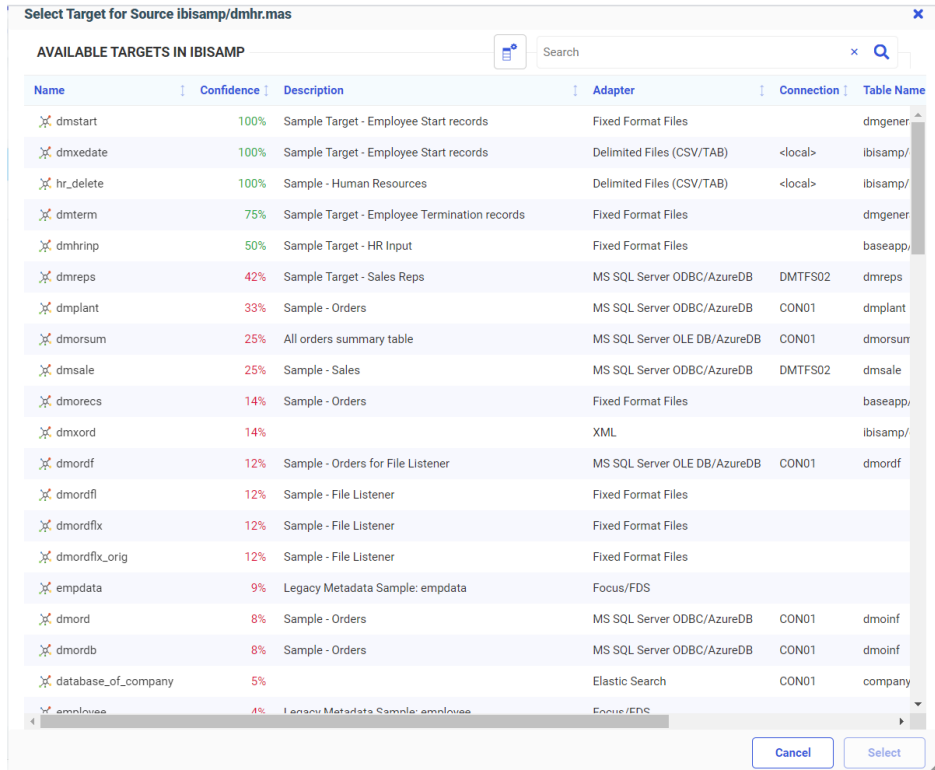


The image shows a dialog box titled "Choose Columns" with a close button (X) in the top right corner. The dialog contains a list of columns with checkboxes next to them. The columns and their selection status are as follows:

Column Name	Selected
Title	<input type="checkbox"/>
Source Name	<input checked="" type="checkbox"/>
Source Adapter	<input type="checkbox"/>
Source Connection	<input type="checkbox"/>
Source Description	<input type="checkbox"/>
Source Table Name	<input type="checkbox"/>
Target Name	<input checked="" type="checkbox"/>
Confidence Level	<input checked="" type="checkbox"/>

At the bottom of the dialog, there are three buttons: "Reset to defaults" (text button), "Cancel" (text button), and "OK" (blue button).

To select an existing target from a list of synonyms, click the target icon. The list of available targets from the Target application directory opens, as shown in the following image.



Generating Data Flows Using a Template Flow

Once you select a target and click *Select*, the folder in the target icon appears with a blue fill color, providing a clear indication that the user made a selection, as shown in the following image.

Generate Flows based on foccache/genflow/template_existing

Configure flow | Template information | Schedule information

Source

Application Folder: ibisamp

Prefix: dm

Existing target

Application Folder: ibisamp

Prefix: _tgt

Flow

Save Location: generatedflow

Process Flow Name: template_existing_g

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name
database_of_company	database_of_company_tgt
dmcomp	comp_tgt
<input checked="" type="checkbox"/> dmhr	dmhrinp
dmhrinp	hrinp_tgt
dminv	inv_tgt
dminv_log	inv_log_tgt
dminva	inva_tgt
dminvd	invd_tgt
dminvf	invf_tgt


Cancel | Generate & Save

- Once you have made all of your changes, click the *Verify* button to check that your selections are valid, as shown in the following image.

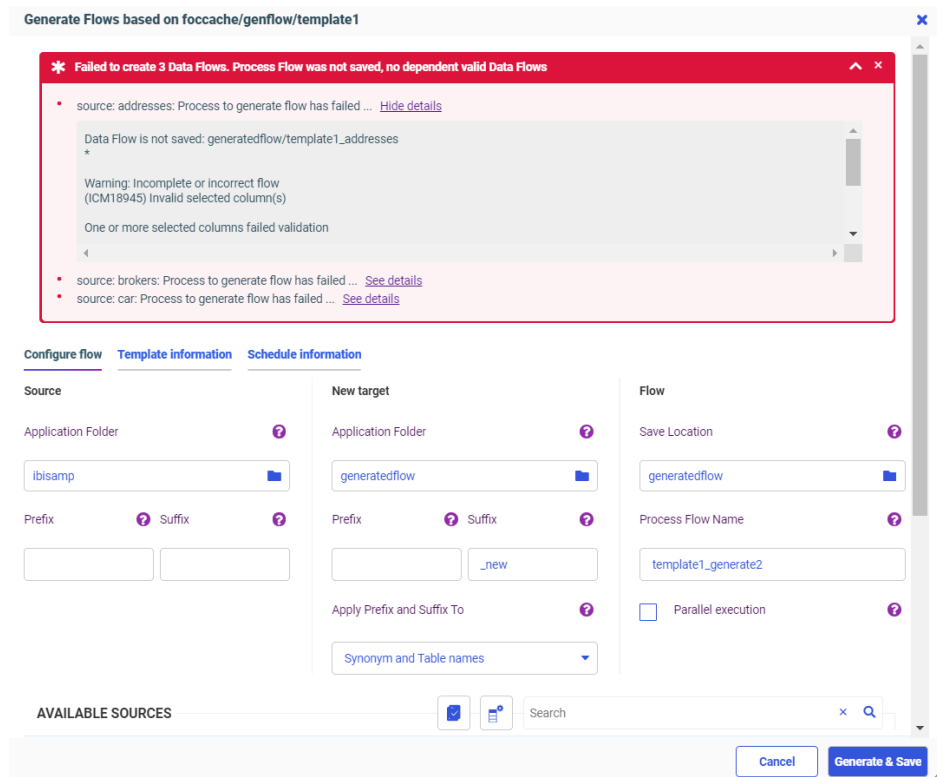
AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
-------------	-------------	-------------------

If there were any problems with your choices, warning messages will display in orange and error messages will display in red at the top of the window, identifying the issues. You can click the down arrow at the top of the widget to see the list of messages. If a message is long, a *See details* link appears, as shown in the following image.



If you click a *See details* link, the Message box expands to show the details. If the details are long, scroll bars enable you to scroll through all of the details, as shown in the following image, in which the sources selected are not valid for the transformations in the flow.



You can click *Hide details* to close the details box.

Generating Data Flows Using a Template Flow

In addition, if rows in the table at the bottom of the screen have validation errors, those rows are highlighted, as shown in the following image, in which the source and target folders and names are the same.

Generate Flows based on foccache/genflow/template01

* Source and target must have different application/synonym names

Configure flow | Template information | Schedule information

Source

Application Folder: ibisamp

Prefix: Suffix:

New target

Application Folder: ibisamp

Prefix: Suffix:

Apply Prefix and Suffix To: Target Name and Target Table Name

Flow

Save Location: generatedflow

Process Flow Name: template01_g

Parallel execution:

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
<input type="checkbox"/> database_of_company	database_of_company	companydatabase/_search
<input checked="" type="checkbox"/> dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/> dmhr	dmhr	dmhr
<input type="checkbox"/> dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/> dminv	dminv	dminv
<input type="checkbox"/> dminv_log	dminv_log	marina.dbo.dminv

Cancel Generate & Save

Note that the *Generate & Save* button is still available, so you can save the erroneous flow.

However, you should correct the errors and click the *Verify* button to validate your changes.

If your choices are valid, a successful verification message displays in green, as shown in the following image.

The screenshot shows a window titled "Generate Flows based on templateflow/template1". At the top, a green banner displays a checkmark and the text "Saved 3 Data Flows, and a Process Flow". Below this, there are three tabs: "Configure flow" (selected), "Template information", and "Schedule information".

The "Configure flow" tab is divided into three sections: "Source", "New target", and "Flow".

- Source:** Application Folder is "ibisamp", Prefix is "dm", and Suffix is empty.
- New target:** Application Folder is "generatedflow", Prefix is empty, Suffix is ".tgt", and "Apply Prefix and Suffix To" is set to "Synonym name".
- Flow:** Save Location is "generatedflow", Process Flow Name is "template2_generate", and "Parallel execution" is unchecked.

At the bottom, there is a table of "AVAILABLE SOURCES" with columns for Source Name, Target Name, and Target Table Name. The table lists several sources, with "dmcomp", "dmhr", and "dminv" selected (checked).

Source Name	Target Name	Target Table Name
database_of_company	database_of_company	database_of_company
<input checked="" type="checkbox"/> dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/> dmhr	dmhr	dmhr
dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/> dminv	dminv	dminv
dminv_log	dminv_log	dminv_log

Buttons for "Cancel" and "Generate & Save" are located at the bottom right of the window.

10. If there were problems identified, correct them.

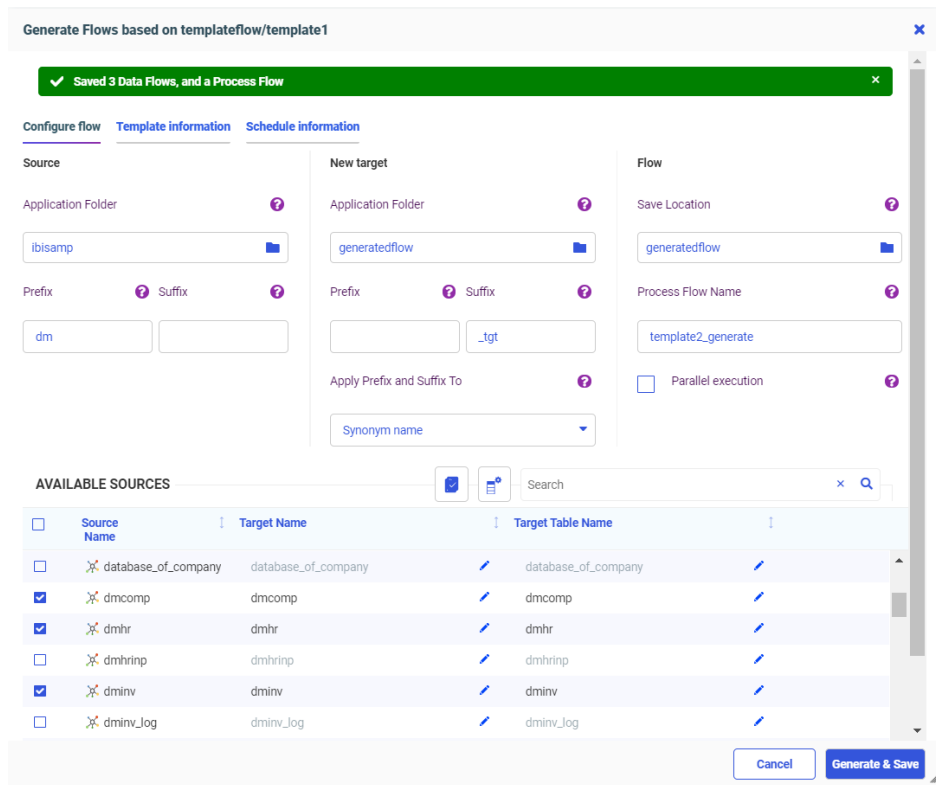
Examples of choices that could prevent verification are:

- The application folders for source and target are the same, with the source names matching the target names. The application folders must be different, unless you change the target names.
- The target synonym name is not valid. It should not exceed 64 characters, and the general guidelines for a synonym name should be followed.
- The target adapter is not available. It must be available (not removed accidentally) at the time when *Generate Flows* is in progress.

- ❑ Transformations you made in the template data flow are not valid for the generated flows. Any transformations you made in the template data flow must be valid for the generated flows. If the source tables do not have the columns necessary to apply the transformations, the flows will not be successfully verified.

11. Once the verification is successful, click either *Generate & Save* or *Generate & Submit*, depending on your selection when you started the Generate Flows process.

Some additional verification processes will be run. If problems occur, messages will display in red at the top of the window, identifying the problems. If no problems occur, a message will display in green, indicating that the save was successful and identifying the objects that were generated, as shown in the following image.







If additional problems were identified, correct them.

If you selected *Generate & Submit* when you started the Generate Flows process, the flows will be submitted, but not saved. Only the target synonyms and target tables will be saved.

12. Close the *Generate Flows* window by clicking *Cancel*.

13. If you selected *Generate & Save*, the generated data flows and the process flow are generated in the flow application directory you selected, as shown in the following image.

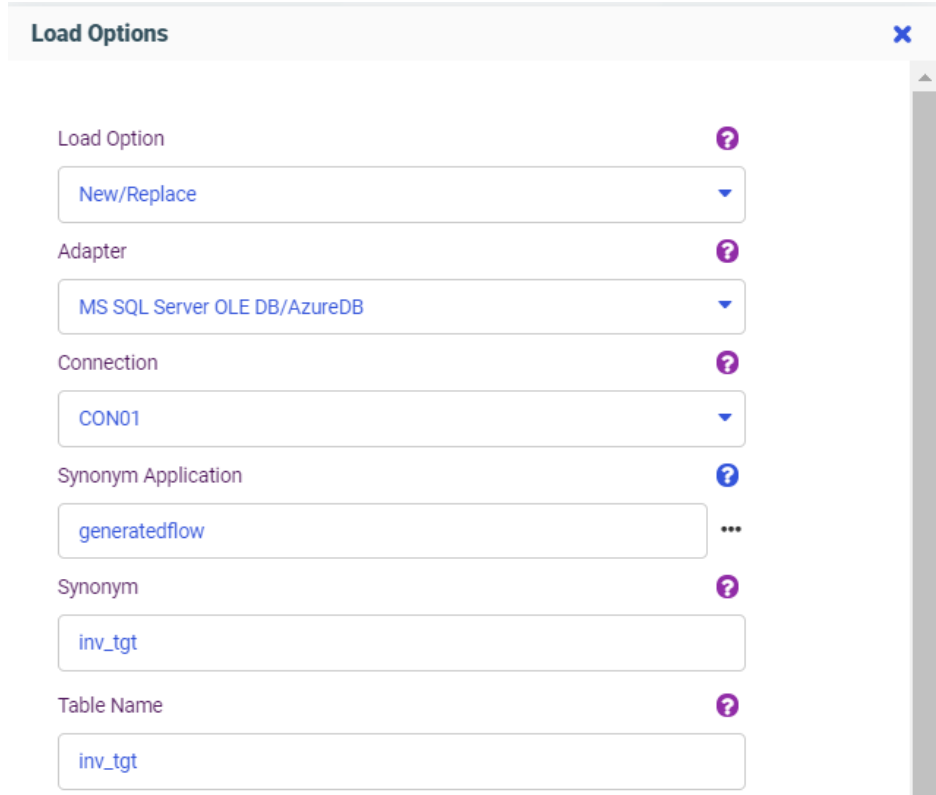
Application Directories/Files ↓	Size ↓	Date Modified ↓
 template1_dmcomp	2197	2020/07/07 10.42.47
 template1_dmhr	2979	2020/07/07 10.42.47
 template1_dminv	2299	2020/07/07 10.42.47
 template1_generate	1698	2020/07/07 10.42.47

The flow named `template1_generate` is the generated process flow. The other flows are the generated data flows, one for each source synonym (identified by having the source synonym name at the end of the flow name).

The following image shows the data flow generated for the `dminv` source synonym. The target synonym name was generated by removing the characters `dm` from the start of the source synonym name and adding the characters `_tgt` to the end of the target synonym name, as specified on the Generate Flows window.




If you selected *Synonym and Table names* from the *Apply Prefix and Suffix To* drop-down list, the prefix and suffix transformations will be applied to the target table name as well as the synonym name, as shown in the following image.




The data flows are generated with the same load options as the template flow, but, you can now open any of the flows and edit them as needed.

To open the process flow, you will have to toggle the flow view by clicking the *Toggle Flow View* button at the top of the window, as shown in the following image.



The following image shows the generated process flow.



14. If you selected *Submit & Save*, once you have finished your edits to the flows, if any, you can submit the generated process flow by right-clicking it and selecting *Submit*.

The targets are generated in the New Target application folder (which, in this example, is the same as the Flow application folder), as shown in the following image.

Application Directories/Files ↓	Size ↓	Date Modified ↓
template1_dmcomp	2197	2020/07/07 10.58.56
template1_dmhr	2979	2020/07/07 10.58.56
template1_dminv	2299	2020/07/07 10.58.56
template1_generate	1698	2020/07/07 10.58.56
hr_tgt	3485	2020/07/07 10.59.15
inv_tgt	985	2020/07/07 10.59.16

Controlling Format Conversion Error Processing in a Data Flow

A format conversion error occurs when a field described with a numeric USAGE format contains character values. In prior releases, format conversion errors in a Data Flow were displayed, but not counted or logged. Records with errors were discarded.

Two parameters can be used to control format conversion error processing in a Data Flow.

- ❑ CNVERR, which instructs the data engine to replace values that cause data truncation or conversion errors with the default value (null if the column is nullable, zero otherwise).

The default value, REJECT, rejects the record. The value IGNORE replaces the values causing errors with default values, and continues processing the record.

- ❑ CNVERRLIMIT, which sets a limit, up to 999999, for the number of records with warnings about data conversion errors.

The default value (0) allows unlimited warning messages. When the limit is reached, processing is terminated and the system variable FOCERRNUM is set to 1346.

The following are the warning messages counted for this limit.

```
(FOC1277) RECORD %1. FIELD %2 HAS FORMAT ERROR, VALUE%3: %4
```

```
(FOC1130) FORMAT CONVERSION ERROR FIELD/KEY %1%2%3%4
```

An error occurred attempting to convert the specified field from its ACTUAL to USAGE format. Either the ACTUAL format is described incorrectly or the data may be invalid. Check the MFD.

```
(FOC1346) %1%2%3%4
```

This message provides the record number and position in this record when data conversion interruption occurs.

Note: These settings apply when the source is a delimited or Excel file.

These parameters are available in the Reporting Server browser interface and in the Source Properties for a Data Flow in the Data Migrator desktop interface.

- ❑ In the Reporting Server browser interface, navigate to the Workspace page. Click *Settings*, point to *FOCUS Sets and Info*, and click *Core Engine Settings*.

Select *Data Utility Service*, and click *Next*. The Data Utility Service page opens, as shown in the following image.


The screenshot shows the 'Data Utility Service' configuration page. It contains several settings, each with a dropdown menu and a description:

- ALPHACLEAN?: OFF (Conversion behavior for external ACTUAL= low value hexadecimal 0.)
- NODATA?: . (Character string display to indicate missing data in a report.)
- OVERFLOWCHAR?: (Character to display during numeric overflow conditions, * (asterisk) by default.)
- EXTENDNUM?: AUTO (Extend numeric usages to avoid format overflow conditions.)
- DISPLAYROUND?: ON (To avoid .499 problem, add ULP (unit of least precision) to display doubles)
- NEG-ZERO?: OFF (Display negative values that round to zero as minus zero.)
- VMSFLOATCONV?: IEEE (OpenVMS Float conversion from G_FLOAT to IEEE for external sources.)
- VMS_TIMESTAMP?: INTERNAL (OpenVMS Timestamp conversion between VMS format and internal format for external sources.)
- CNVERR?: REJECT (Data conversion / truncation error processing.)**
- CNVERRLIMIT?: 0 (Limit number of records with warnings about data conversion error. 0 - unlimited.)**
- PHONETIC_ALGORITHM?: METAPHONE (Phonetic algorithm for indexing of strings in function PHONETIC.)
- ALLOWCVTERR?: OFF (Date conversion behavior for external ACTUAL= to USAGE= data errors.)

At the bottom right, there are three buttons: 'Cancel', 'Save', and 'Back'.

- ❑ In the Data Migrator desktop interface, right-click a delimited, fixed-format, or Excel source, and click *Properties*.

The Properties panel opens, as shown in the following image.



Migrating Data Flows

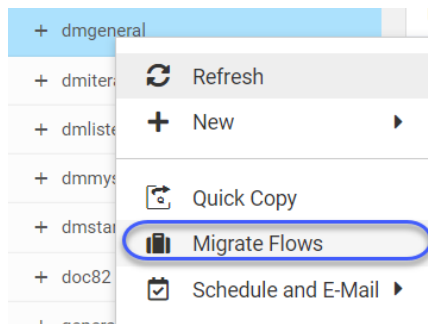
When you create a Data Flow, an SQL request is generated that retrieves data from Source databases, applies transformations, if necessary, and loads the data to the Target databases.

In new releases of Data Migrator and the server, improvements are made, when possible, to the code that is generated for the Flows. The new code uses new back-end algorithms for reading and loading the data, allowing improved performance.

If a Data Flow created in a prior release no longer runs, or seems to run too slowly, you can migrate the flow to the current release. This process will update the code for the generated Flow to the latest syntax, and take advantage of performance and optimization improvements available in the current release.

Procedure: How to Migrate Data Flows

1. Right-click an application directory, and click *Migrate Flows*, as shown in the following image.



The *Migrate Flows* window opens, as shown in the following image.

Migrate Flows

From

Application Folder

To

Application Folder

Options

Optimize Load

FLOWS TO MIGRATE IN DMGENERAL

<input type="checkbox"/>	Name	Description	Server Release	Client Release	Last Modified
<input type="checkbox"/>	create_dmordf_dmordfl		R729999D.300968	R729999D.300968	2020/06/18 09.23.22
<input type="checkbox"/>	dmrcode	Decode function used inline and from file			2020/06/09 11.40.51
<input type="checkbox"/>	dmdemo	Join with source and target transformations, sql calculations			2020/06/09 11.40.51
<input type="checkbox"/>	dmduplg	Record logging for rejected rows			2020/06/09 11.40.51
<input type="checkbox"/>	dmflow01	Sales Counts Initial Load	R728207D.1925		2020/10/29 11.56.19
<input type="checkbox"/>	dmikups	Lookup in two tables			2020/06/09 11.40.51
<input type="checkbox"/>	dmordxml	Target (XML) described by dmxord.xsd			2020/06/09 11.40.51
<input type="checkbox"/>	dmorecs	Rectype type used to create one file with two record types			2020/06/09 11.40.51
<input type="checkbox"/>	dmpeople	Use UNION to combine two select statements			2020/06/09 11.40.51


Migrate

2. Optionally, change the application directories for the migration process and deselect the Optimize Load option.

Initially, both the *From* and *To* application folders are set to the application directory from which you initiated *Migrate Flows*.

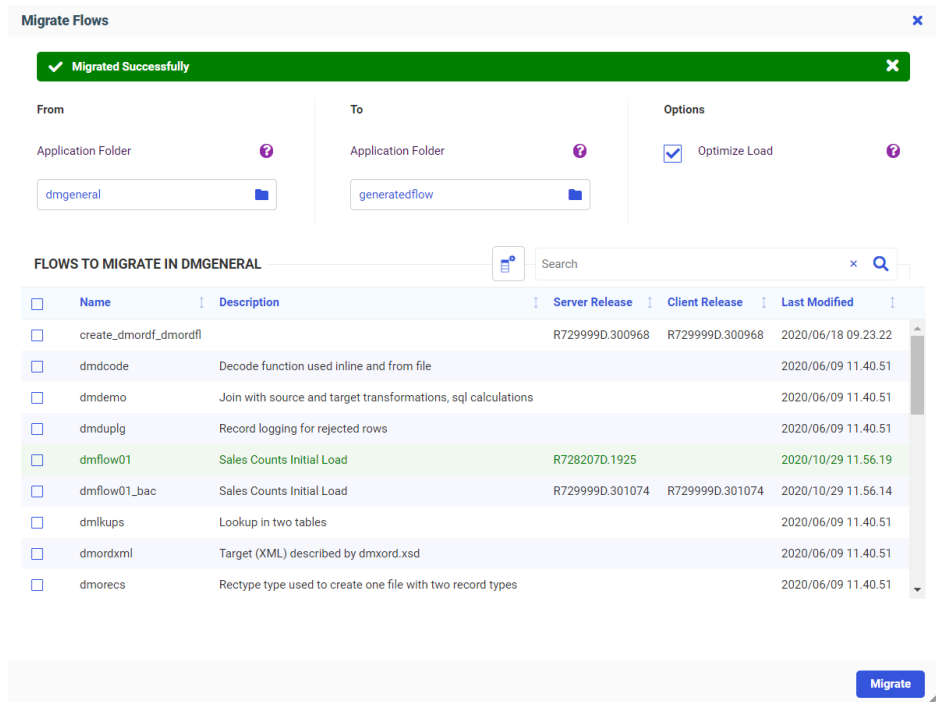
- Click in the *Application Folder* entry field to open the *Select Application Folder* dialog box to change one or both of the folders.
 - By default, *Optimize Load* is selected. If you want to disable this feature, deselect the check box.
3. Select the check boxes next to the flows you want to migrate, and click *Migrate*.

If the *From* and *To* folders are the same, the original versions of the selected flows will be saved with the suffix `_bac` appended to the flow name. The migrated flows will have the original names. The following message displays to inform you of this.



If you receive this message, either click *OK* to migrate the flows and save the original versions with names that have the suffix `_bac`, or click *Cancel* to cancel the migration.

The Migrate Flows window has either a green banner and success message at the top, if the flows were migrated successfully, or a red banner with error messages, if they were not. A successful migration is shown in the following image.



The flows on the list that were selected for migration are also highlighted.

4. Close the Migrate Flows window by clicking the X at the top right, if you are using the Reporting Server browser interface.
 - If the *From* and *To* application folders were the same, the original flows (with the *_bac* suffix) and migrated flows will appear in the application folder.
 - If the *From* and *To* folders were different, the original flows will be in the *From* folder, with their original names, and the migrated flows will be generated in the *To* folder, with the same names.

Viewing Checkpoint Files Generated by a Change Data Capture (CDC) Process

When you run a flow, a Checkpoint file is generated or, if it already exists, is updated. The Checkpoint file is in a proprietary internal binary format, so users cannot read it to see what happened during the processing. A viewer for checkpoint files is available that will show the checkpoint file contents in readable JSON format.

Note: The contents of a checkpoint file are specific to the DBMS being used.

In order for CDC to be enabled, the tables used must be log-enabled. The Data Flow is then run using the Log file as the source and Change Data Capture as the load option.

1. Create a synonym for the Table Log Records, as shown in the following image. You can create a new application folder to contain the log synonym and log file.

Note: This example uses Db2 tables on a server running on IBM i.

Create Synonym for DB2/DB2 Warehouse (CON1)

Object Type Library Supply value to avoid system-wide search, sample: S999

Table Name

Miscellaneous Parameters

Customize data type mappings

Create: Cluster Synonym Cluster Join for Star Schema Base Synonym

Application Prefix Suffix

SYNONYM CANDIDATES Row Limit: 50 Search

Select	Default Synonym Name	Table Name	Library/Schema	Type
<input type="radio"/>	c1_tmix91b	C1_TMIX91B	DD999	Log Records
<input type="radio"/>	c1_tmix91c	C1_TMIX91C	DD999	Log Records
<input type="radio"/>	c1_tmix91d	C1_TMIX91D	DD999	Log Records
<input type="radio"/>	c1_tmix91e	C1_TMIX91E	DD999	Log Records
<input type="radio"/>	c1_tmix91m	C1_TMIX91M	DD999	Log Records
<input type="radio"/>	c1_trdddc1	C1_TRDDCDC1	DD999	Log Records
<input type="radio"/>	c1_trdddc2	C1_TRDDCDC2	DD999	Log Records
<input type="radio"/>	c1_trdxmla	C1_TRDXMLA	DD999	Log Records

2. Select the table for which you want to create a synonym. If necessary, add a prefix or suffix that will identify this as a synonym for a log table, optionally change the default synonym name, and click *Add*.

The Synonym Editor opens.

3. Make any changes that are necessary, and click the right arrow on the ribbon to save the synonym.

The synonym contains variables that control the CDC processing. You can edit them in the Synonym Editor or in the Text Editor. To open the synonym in the Text Editor, right-click the synonym on the Resources tree, point to *Metadata Management*, and click *Edit as Text*.

The synonym opens as text, as shown in the following image.

```
FILENAME=C1_TRDCDC1_LOG, SUFFIX=DB2      , $
VARIABLE NAME=&&CDC_POLLING,
  PROMPT='DBMS log polling interval (in seconds) for LUW',
  USAGE=I5, DEFAULT='1', $
VARIABLE NAME=&&CDC_TIMEOUT,
  PROMPT='Timeout interval (in seconds) to listen DBMS log for LUW',
  USAGE=I5, DEFAULT='1', $
VARIABLE NAME=&&CDC_START,
  PROMPT='Starting point in reading log records',
  DEFAULT='CHKPT', ACCEPT=
  DECODE(<'CHKPT - After last LUW retained in checkpoint file',
  'CHKPT'>, <'CUR_TRAN - First LUW after DBMS log reading started',
  'CUR_TRAN'>, <'CUR_LOG - First available LUW in DBMS log',
  'CUR_LOG'>), $
VARIABLE NAME=&&CDC_MAXLUWS,
  PROMPT='Maximum number of LUWs processed in the request',
  USAGE=I10, DEFAULT='1', $
VARIABLE NAME=&&CDC_CHKPT_SAVE,
  PROMPT='Retain last processed LUW in checkpoint file',
  DEFAULT='YES', ACCEPT=
  DECODE(<'YES - Retain last processed LUW in checkpoint file', 'YES'>,
  <'NO - Don''t retain LUWs in checkpoint file', 'NO'>), $
VARIABLE NAME=&&CDC_CHKPT_FILE,
  PROMPT='Checkpoint file location', USAGE=A99, DEFAULT=' ', $
VARIABLE NAME=&&CDC_LOG_NAME,
  PROMPT='DBMS log file name', DEFAULT=' ', $
VARIABLE NAME=&&CDC_LOG_LOCATION,
  PROMPT='DBMS log file location', DEFAULT=' ', $
VARIABLE NAME=&&CDC_COMMIT_MODE,
  PROMPT='Transactions Commit Mode',
  DEFAULT=' ', ACCEPT=
  DECODE(<'ON - Transactions Committed', 'ON'>,
  <'OFF - Transactions Autocommit', 'OFF'>), $
```

```

SEGMENT=C1_TRDCDC1_LOG, SEGTYPE=S0, $
  FIELDNAME=CDC_OPER, ALIAS=__C1, USAGE=A1, ACTUAL=A1, $
  FIELDNAME=CDC_TID, ALIAS=__C2, USAGE=A32, ACTUAL=A32, $
  FIELDNAME=CDC_TIMES, ALIAS=__C3, USAGE=A26, ACTUAL=A26, $
  FIELDNAME=FA01INT, ALIAS=FA01INT, USAGE=I11, ACTUAL=I4, $
  FIELDNAME=FA02INT, ALIAS=FA02INT, USAGE=I11, ACTUAL=I4, $
  MISSING=ON, $
  FIELDNAME=FA03REAL, ALIAS=FA03REAL, USAGE=D20.2, ACTUAL=D8, $
  MISSING=ON, $
  FIELDNAME=FA04NUMERIC16X4, ALIAS=FA04NUMERIC16X4, $
  USAGE=P19.4, ACTUAL=P9, MISSING=ON, $
  FIELDNAME=FA05NUMERIC, ALIAS=FA05NUMERIC, USAGE=P16, ACTUAL=P8, $
  MISSING=ON, $
  FIELDNAME=FA06CHAR_5, ALIAS=FA06CHAR_5, USAGE=A5, ACTUAL=A5, $
  MISSING=ON, $
  FIELDNAME=FA07VARCHAR_5, ALIAS=FA07VARCHAR_5, USAGE=A5V, ACTUAL=A5V, $
  MISSING=ON, $
  FIELDNAME=FA08CHAR_10, ALIAS=FA08CHAR_10, $
  USAGE=A10, ACTUAL=A10, MISSING=ON, $
  FIELDNAME=FA09VARCHAR_10, ALIAS=FA09VARCHAR_10, $
  USAGE=A10V, ACTUAL=A10V, MISSING=ON, $
  FIELDNAME=FA10DATE, ALIAS=FA10DATE, USAGE=YYMD, ACTUAL=DATE, $
  MISSING=ON, $
  FIELDNAME=FA11TSTAMP, ALIAS=FA11TSTAMP, USAGE=HYMDm, ACTUAL=HYMDm, $
  MISSING=ON, $
  FIELDNAME=FA12TIME, ALIAS=FA12TIME, USAGE=HHIS, ACTUAL=HHIS, $
  MISSING=ON, $

```

4. Edit the variable values, as follows:

- &&CDC_MAXLUWS.** Optionally, change this to zero (0) in order to add all LUWs to the checkpoint file.
- &&CDC_CHKPT_SAVE.** Make sure this is set to YES, in order to retain the last processed checkpoint in the checkpoint file.
- &&CDC_CHKPT_FILE.** Optionally, add the checkpoint file name. The extension must be .chp.


For example:

```

VARIABLE NAME=&&CDC_MAXLUWS,
  PROMPT='Maximum number of LUWs processed in the request',
  USAGE=I10, DEFAULT='0', $
VARIABLE NAME=&&CDC_CHKPT_SAVE,
  PROMPT='Retain last processed LUW in checkpoint file',
  DEFAULT='YES', ACCEPT=
  DECODE(<'YES - Retain last processed LUW in checkpoint file',
  'YES'>,
  <'NO - Don't retain LUWs in checkpoint file', 'NO'>), $
VARIABLE NAME=&&CDC_CHKPT_FILE,
  PROMPT='Checkpoint file location',
  USAGE=A99, DEFAULT='mycheck.chp', $

```

5. Run the flow with the log table as the source and your existing target, as shown in the following image.



Change the Load Option to Change Data Capture, as shown in the following image.

Load Options [X]

Load Option [Change Data Capture] ?

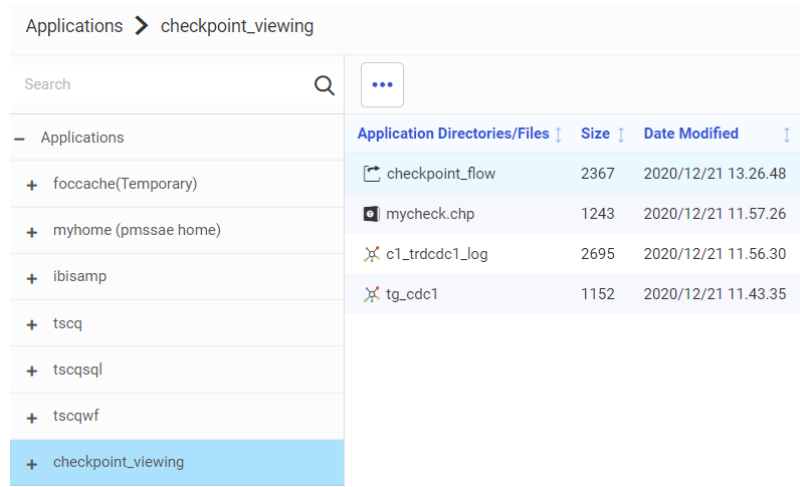
Prior to Load Option [None] ?

Select Target Synonym [checkpoint_viewing/tg_cdc1.mas] ?

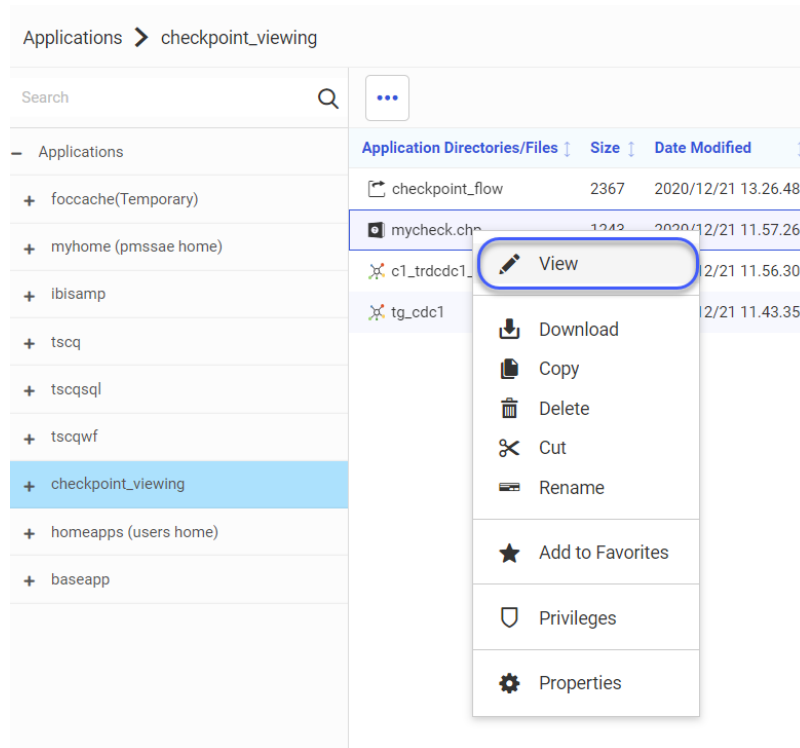
Control column [CDC_OPER] ?

[Cancel] [OK]

The log file is generated in the application folder, as shown in the following image.



- To view the checkpoint file, either double-click the file or right-click the file and click *View*, as shown in the following image.



The checkpoint file is displayed in JSON format, as shown in the following image.

```
{
  "checkpoint": [
    {
      "name": "file",
      "value": "checkpoint_viewing/mycheck.chp",
      "description": "check point file full name"
    },
    {
      "name": "engine",
      "value": "DB2",
      "description": "data access adapter"
    },
    {
      "name": "catalog",
      "value": "*LOCAL",
      "description": "catalog / database name"
    },
    {
      "name": "table",
      "value": "DD999.C1_TRDCDC1",
      "description": "table name"
    },
    {
      "name": "chkp_id",
      "value": "###00000000000000000000",
      "description": "checkpoint ID / min entry LSN"
    },
    {
      "name": "commit_id",
      "value": "QSQRN0002##DD999      ##00000000000000199545",
      "description": "commit ID/LSN"
    },
    {
      "name": "tran_id",
      "value": "00000000000000199533",
      "description": "current transaction ID/LSN"
    },
    {
      "name": "timestamp",
      "value": "2020-12-21 06:27:58.630656000",
      "description": "transaction timestamp"
    }
  ]
}
```

Stored Procedures

Procedures are reusable components of application logic which can include reporting or extract syntax, and are usually written in the native 4GL language of the server (TABLE). The native language of the server also includes syntax for controlling the execution of events within the request, and is known as Dialogue Manager. Procedures that contain TABLE and Dialogue Manager commands are called FOCEXECs.

Procedures may also be written in compiled and linked 3GL languages, such as C, COBOL, Fortran, Assembler, RPG, and others, and then can be called from the native 4GL language of the server.

This chapter discusses use of tools to create 4GL procedures plus an introduction to the general use of the Dialogue Manager language. The *TIBCO WebFOCUS® Stored Procedure and Subroutine Reference for 3GL Languages* manual discusses the creation of procedures using 3GL languages.

In this chapter:

- Working With Stored Procedures
- Creating a Stored Procedure
- Using the Reporting Server Browser Interface Search
- Sending Stored Procedure Report Output to Server-Side Printers
- Editing a Stored Procedure
- Running a Stored Procedure
- Debugging a Stored Procedure
- Running a Procedure With Advanced Options
- Scheduling a Procedure
- Sending an E-mail Notification for a Procedure
- Procedure Run-Time Processing Essentials
- Platform-Specific Commands and Features
- Stress Testing a Procedure
- Running Impact Analysis Reports

Working With Stored Procedures

Procedures are available from the navigation pane of the Application page. You can:

- Create and edit stored procedures.

- Run procedures at the present time, at a later time, or under stress conditions (Reporting Server browser interface only).
- Run a variety of reports that monitor many aspects of procedure performance.

Creating a Stored Procedure

You can code a stored procedure in the text editor feature that is easily accessible from the Reporting Server browser interface and the Data Migrator desktop interface.

The contents of a procedure can vary greatly. Reports can be created that go to a user's browser, a configured printer, or held as a data extract in various formats. The reporting and extract language is called TABLE. Additionally, syntax is available to parameterize a request and control the flow of execution in a request. Parameters (called amper variables) and flow of control within a procedure are part of the Dialogue Manager language.

For complete information about TABLE, see the *Creating Reports With TIBCO WebFOCUS® Language* manual. For complete information about Dialogue Manager commands, see the *TIBCO WebFOCUS® Developing Reporting Applications* manual. While the entire syntax of these features is not within the scope of this manual, it is necessary to provide information about basic run-time essentials. For information, see [Procedure Run-Time Processing Essentials](#) on page 495.

The following is a small representative example that embodies reporting (TABLE), parameters (Amper Variables) and execution flow (Dialogue Manager):

```
-DEFAULT &COUNTRY = ENGLAND
-IF &COUNTRY EQ '' THEN GOTO NOPARM ;
TABLE FILE CAR
PRINT COUNTRY CAR
WHERE COUNTRY = '&COUNTRY' ;
END
-EXIT
-NOPARM
TABLE FILE CAR
PRINT COUNTRY CAR
END
```

Procedure: How to Create a Stored Procedure

You can create procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Right-click an application folder, point to *New*, then click *Procedure*.

A text editing window opens.

3. Type the procedure code, which may be either FOCUS syntax or SQL Passthru syntax. For example:

```
SQL DB2 SELECT x FROM y ;
END
```

When using a pure editor environment, you are responsible for the accuracy of the syntax you enter. However, any errors that are generated will be displayed in the right pane, as well as in the Session Log, which you can access from the *My Console* option on the menu bar.

4. Click the *Run* or *Save As* icon.

When you run the procedure, the output is displayed in the right pane.

When you Save the procedure, the Save As dialog box opens.

5. Select an application folder.
6. Enter a name in the *File Name* field.
7. Click *OK*.


The procedure is added to the application folder in the resources tree.

Using the Reporting Server Browser Interface Search


The search is available on all Reporting Server browser interface pages that display a list of files. For example, the search is available on the Synonym Candidates pages, the Manage Files page, and the Application Path Configuration page.

Clicking the magnifying glass icon opens a Find box where you can enter characters. Click the magnifying glass next to the text box to implement the search.

You can also click the plus sign (+) to open a Replace entry field, as shown in the following image.



When you enter characters in the *Search for* text box, the first occurrence is highlighted. To highlight all occurrences, click *All*.



All matching characters on the list will be highlighted. Depending on where the search is implemented, forward and backward arrows may be available to go from one found instance to the next or previous instance.

There are buttons for the following options:

- Regular Expression Search.** You can enter a search pattern in the *Search for* text box.
- Case Sensitive Search.** Check the box to do a case-sensitive search.
- Whole Word Search.** Finds strings that contain only the specified characters.
- Search in Selection.** Finds strings within the selected text.

If you enter a string in the Replace text box, you can click *Replace* to replace the current instance and move to the next matching string, or *All* to replace all occurrences.

There is a *Select All* button next to the Search button.

Sending Stored Procedure Report Output to Server-Side Printers

A report is spooled to print using a combination of commands. You must issue a FILEDEF (or ALLOCATE on z/OS) command to point to a specific printer, the OFFLINE command to redirect the report request output to the device, the report request itself, and the OFFLINE CLOSE command to close the spooling session (as multiple requests may be spooled as part of a single print output). For details, see [OFFLINE Printing](#) on page 561.

Editing a Stored Procedure

Once you have created a procedure, you can open the text editor to view or revise it. The text editor provides sophisticated search and replace functions, code folding, and autocompletion options.

You can open a Master or Access File in the text editor by right-clicking the synonym on the file list and clicking *Edit as Text* (to edit the Master File) or *Edit Access File as Text* (to edit the Access File).

Some types of files open in ready-only mode in the text editor, such as when you open a Master File for viewing in Data Assist. In those situations, some text editor options are not available.

Tip: In addition to editing procedure code, you can copy, delete, or move a procedure component using options on the *Applications* resources tree or file list.

You can edit procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.

3. Right-click a procedure, and select *Open*.

The procedure opens in the text editor.

Tip: Double-clicking the procedure will also open the text editor.

4. When you have made your changes, click the *Save*, *Save As*, or *Run* icon.

Text Editor Menu Bar Functions

The text editor menu bar is shown in the following image.



The following functions are available.

- Save.** Saves the current text under the same name.
- Save As.** Saves the current text under a new name.
- Run.** Runs the current text, if it is a procedure.
- Tab Navigation.** Enables cycling through the interface using the Tab key.
- Find.** Opens the search function.
- Selection folding.** Folds unselected lines.
- Uppercase.** Changes selected text, or the token with the cursor, to uppercase. A token is an unbroken string of alphanumeric characters and underscores.
- Lowercase.** Changes selected text, or the token with the cursor, to lowercase. A token is an unbroken string of alphanumeric characters and underscores.
- Comment/Uncomment.** Toggles comments on and off for the current line or selected lines.
Note: If a line is already commented, and you click the *Comment/Uncomment* button, new comment characters are added in addition to the comment characters that already existed. In this way, if you click the button again, the line goes back to its original state (commented). This prevents you from unintentionally uncommenting a line that was meant to be a comment.
- Undo.** Undoes any number of edits.
- Redo.** Redoes any number of edits.
- Select All.** Selects all of the text in the text editor.

- Options.** Opens the text editor options dialog box.


Text Editor Color Coding

Certain syntax components in FOCEXECs and Master Files are color-coded in the text editor.

- Strings are displayed in dark green.
- Comments are displayed in light green.
- FOCEXEC keywords are displayed in blue.
- Numbers are displayed in blue.
- Expression keywords are displayed in gray.
- Amper variables are displayed in light brown.
- SQL keywords are displayed in purple.

Searching for Text or a Regular Expression

Clicking *Find* (magnifying glass) on the menu bar opens the search panel. Clicking the plus sign (+) under the Search for text box opens the Replace section, as shown in the following image.



The following Find options are available.

- Search for.** Enter text in the input field. If you click the Regular Expression button (.*) at the bottom, enter a regular expression. If text is selected when you click Find, the text will automatically be entered in the *Search for* input field. When there is text in the *Search for* field, clicking the x next to the term deletes it.

Note: The bottom row of the Find panel identifies which instance has the focus and the total number of matches found, in the form *n* of *m*.

- Find previous (<)/Find next (>).** Enables you to go from instance to instance of matching text.
- All.** Click *All* to highlight all matching instances and place a cursor on each of them for bulk deletion or replacement.

- Replace with.** Enter replacement text in the input field.
 - Replace.** Click *Replace* to replace the current instance.
 - All.** Click *All* to replace all matching instances.
- RegExp Search (.*)**. Click *.** to enter a JavaScript regular expression in the *Search for* input field.
- CaseSensitive Search (aA)**. Click *aA* to toggle between case-sensitive and case-insensitive search.
- Whole Word Search (\b)**. Click *\b* to toggle between accepting matches within words and whole words only. A word is an unbroken string of alphanumeric characters and underscores.
- Search in Selection (S)**. Searches within the current selection.

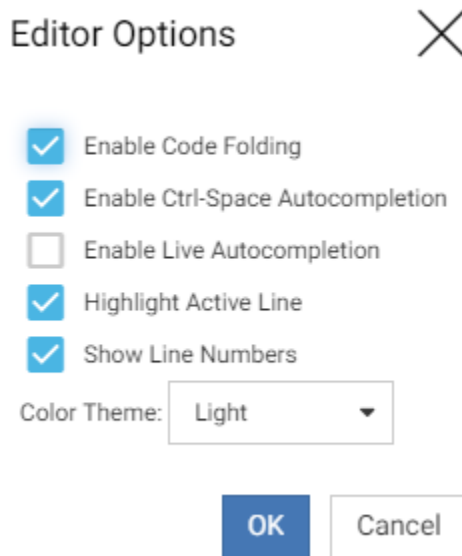
The following image shows a regular expression search for the colon, left brace, and left bracket characters, with the focus on instance 2 of 40.



To close the Find function, click *x* at the right of the Search for row.

Configuring Text Editor Options

Pointing to the down arrow at the right end of the text editor menu bar and clicking *Text Editor Options* opens the Editor Options dialog box, as shown in the following image.



The following options are available. Click the check box next to an option to toggle the selection on and off.

- Enable Code Folding.** Code folding places down arrows next to certain commands and characters that flow onto multiple lines, so that you can compress them to hide their content. When code is folded, the arrow changes to a right-facing arrow that you can click to unfold the syntax. This option is selected by default.
- Enable Ctrl-Space Autocompletion.** Provides a list of keywords available in the current context when you click the Ctrl key and the space bar together. This option is selected by default.
- Enable Live Autocompletion.** Provides a list of keywords available in the current context whenever you place the cursor in the text editor at a point where a keyword is available and you start typing. This option is unselected by default.
- Highlight Active Line.** Highlights the line with the cursor in yellow. This option is selected by default.

- Show Line Numbers.** Displays line numbers to the left of each line. This option is selected by default.
- Color Theme.** Changes the color theme for the current session. Choices are Light, which is the default, and Dark, which makes the background dark blue.

When you have finished selecting the options you want, click *OK*.

Using Code Folding

When code folding is enabled, certain commands and syntax components that flow onto multiple lines can be rolled up so their content is hidden. Lines that can be folded display a down arrow to the left of the line, as shown in the following image.

```

1 ▾ DEFINE FILE WF_RETAIL_LITE
2 ▾ COGS_NEW = COGS_US + (
3           COGS_US * .05);
4 END
5 ▾ GRAPH FILE WF_RETAIL_LITE
6 SUM COGS_US GROSS_PROFIT_US REVENUE_US DISCOUNT_US
7 ACROSS PRODUCT_CATEGORY
8 ON GRAPH PCHOLD FORMAT JSCHART
9 ON GRAPH SET LOOKGRAPH VBAR
10 ON GRAPH SET STYLE *
11 *GRAPH_JS
12 ▾ yaxis: {title: {visible: true, color: 'red', font: 'bold 10pt Verdana',
13   text: 'Y-AXIS'},
14   bodyLineStyle: {width: 2, color: 'red'},
15   swapChartSide: true},
16 ▾ y2axis: {title: {visible: true, color: 'red',
17   font: 'bold 10pt Verdana',text: 'Y2-AXIS'},
18   bodyLineStyle: {width: 2, color: 'green'}},
19 ▾ series: [
20   {series: 0, color: 'cyan'},
21   {series: 1, color: 'tan'},
22   {series: 2, color: 'lightblue'},
23   {series: 3, color: 'teal',
24     yAxisAssignment: 2}]
25 *END
26 ENDSTYLE
27 END

```

The following image shows the same content with the DEFINE FILE command folded.

```

1 ▸ DEFINE FILE WF_RETAIL_LITE
5 ▾ GRAPH FILE WF_RETAIL_LITE
6 SUM COGS_US GROSS_PROFIT_US REVENUE_US DISCOUNT_US
7 ACROSS PRODUCT_CATEGORY
8 ON GRAPH PCHOLD FORMAT JSCHART
9 ON GRAPH SET LOOKGRAPH VBAR
10 ON GRAPH SET STYLE *
11 *GRAPH_JS
12 ▾ yaxis: {title: {visible: true, color: 'red', font: 'bold 10pt Verdana',
13 text: 'Y-AXIS'},
14 bodyLineStyle: {width: 2, color: 'red'},
15 swapChartSide: true},
16 ▾ y2axis: {title: {visible: true, color: 'red',
17 font: 'bold 10pt Verdana',text: 'Y2-AXIS'},
18 bodyLineStyle: {width: 2, color: 'green'}},
19 ▾ series: [
20 {series: 0, color: 'cyan'},
21 {series: 1, color: 'tan'},
22 {series: 2, color: 'lightblue'},
23 ▾ {series: 3, color: 'teal',
24 yAxisAssignment: 2}]
25 *END
26 ENDSTYLE
27 END

```

The arrow next to the DEFINE FILE command has changed to a right-facing arrow. Clicking it unfolds the command.

The following commands and characters can be folded.


- DEFINE FILE.
- TABLE FILE.
- GRAPH FILE.
- MATCH FILE.
- MODIFY FILE.
- MAINTAIN FILE.
- SQL.
- Parentheses, braces, and brackets that open on one line and close on another line.
- SEGNAME/SEGMENT in a Master File.
- END in a Master File (this folds the DBA section of the Master File).

Using Custom Folding

Custom folding enables you to choose the type of commands that you want to remain unfolded. This is especially useful when you open the Session Log, which contains all of the commands generated in your session.

To enable custom folding, click the down arrow on the text editor toolbar, and click *Custom Folding...*

The Custom Folding dialog box opens, as shown in the following image.



Check the boxes for the commands that you *do not* want to be folded. For example, DEFINE is checked for the following image that shows a file with comments, SET commands, a DEFINE command, and TABLE request.

```
1  CREATE TABLE  
8  DEFINE FILE retail18205/bv_namespace_off/wf_retail_lite_off ADD  
9  TIME_YEAR_GROUP/16 TITLE "Year,Group"-If TIME_YEAR IN (2008, 2009, 2010, 2011) THEN 0 ELSE IF TIME_YEAR IN (2012, 2013, 2014) THEN 1 ELSE TIME_YEAR;  
10 END  
11  CREATE TABLE
```

If you previously implemented custom folding and want to remove it, click *Remove Custom Folding*. You can also unfold any folded text by clicking it.

Click *OK*.

Using Selection Folding

If you make selections in the editor and click the *Selection Folding* button, all lines that were not selected are hidden so that only the selected lines display. The following image shows selection folding after searching for occurrences of the characters *LITE* in the *WF_RETAIL_LITE* Master File.

```

1
2 FILENAME=WF_RETAIL_LITE, TRANS_FILE=EDAHOME/NLS/dt,
3 REMARKS='WR LITE Fact/Dim Cluster', BV_NAMESPACE=OFF, S
54 * SEGMENT=WF_RETAIL_TIME_SALES, CRFILE=retail18205/dimensions/wf_retail_time_1166, CRSEGMENT=WF_RETAIL_TIME_LITE, CRINCLUDE=ALL, CRJOINT
57 * SEGMENT=WF_RETAIL_TIME_DELIVERED, SEGTYPE=KU, PARENT=WF_RETAIL_SHIPMENTS, CRFILE=retail18205/dimensions/wf_retail_time_1166, CRSEGMENT=WF_RETAIL_TIME_LITE, CRINCLUDE=ALL, CRJOINTYPE=LI
60 * SEGMENT=WF_RETAIL_TIME_SHIPPED, SEGTYPE=KU, PARENT=WF_RETAIL_SHIPMENTS, CRFILE=retail18205/dimensions/wf_retail_time_1166, CRSEGMENT=WF_RETAIL_TIME_LITE, CRINCLUDE=ALL, CRJOINTY

```

This feature is useful when you make multiple non-contiguous selections, so that you can easily view all of your selections. The button toggles selection folding on and off, so that once selections are folded, you can click the button again to unfold them.

Text Editor Autocompletion

Autocompletion opens a list of supported keywords as you enter text in the Text Editor. Autocompletion is context-sensitive. That is, it will only display keywords available in the portion of the syntax that has the focus. If you select a keyword, autocompletion does not parse the syntax to make sure it is correct.

You can enable autocompletion on an as requested basis (Ctrl+Space autocompletion) or whenever you are entering text where a keyword is available (live autocompletion).

The following image shows an autocompletion list within a GRAPH FILE command.

```
1 ▸ DEFINE FILE WF_RETAIL_LITE
5 ▾ GRAPH FILE WF_RETAIL_LITE
6 SUM COGS_US GROSS_PROFIT_US REVENUE_US DISCOUNT_US
7 A| PRODUCT_CATEGORY
8 ACROSS TABLE/GRAPH
9 ADD TABLE/GRAPH
10 AFTER TABLE/GRAPH
11 AND TABLE/GRAPH
12 ▾ AS TABLE/GRAPH
13 AS TABLE/GRAPH
14 ASQ. TABLE/GRAPH
15 AS_SPAN TABLE/GRAPH
16 ▾ AVE. TABLE/GRAPH
17 font: 'bold 10pt verdana',text: 'Y2-AXIS'},
18 bodyLineStyle: {width: 2, color: 'green'}},
19 series: [
20 {series: 0, color: 'cyan'},
21 {series: 1, color: 'tan'},
22 {series: 2, color: 'lightblue'},
23 ▾ {series: 3, color: 'teal',
24 yAxisAssignment: 2}]
25 *END
26 ENDSTYLE
27 END
28
```

Click a keyword on the list to insert it into the syntax.

Using Tab Key Navigation

The text editor Tab key navigation supports Section 508 Accessibility Compliance by enabling users to navigate through the interface using the Tab key on the keyboard.

To implement Tab key navigation, either press the *Esc* key when in a file open in the text editor, or click the Tab Navigation button on the toolbar that is shown in the following image:



Once you have enabled Tab key navigation, the icon on the button becomes hollow in the center, as shown in the following image.



With Tab key navigation enabled, pressing Tab cycles between the features on the page.

- If you press *Tab* from within a file open in the text editor, the focus moves to the URL bar at the top of the web page, then to the actual URL.
- Continuing to press Tab cycles to the buttons on the ribbon, with a change in color intensity identifying which button has the focus.
- After the ribbon buttons, the focus moves to the data tree, with a cursor or outline showing the position of the focus.
- From there, the focus moves to the text editor toolbar, indicated by a rounding of the edges of the button that has the focus.
- Pressing *Enter* selects the option that has the focus.

Comparing or Merging Files

Using the server Compare Files and Merge Files options, you can compare or merge two files in the server text editor.

To compare or merge files, multi-select two files (by clicking one file, holding down the Ctrl key (or the Shift key if the files are contiguous) and clicking the second file), right-click, and select either *Compare Files* or *Merge Files*. Although both can show the differences, only Merge Files enables you to move content between the files. Compare Files is read-only.


Note: If the two files represent synonyms, selecting:

- Compare Files* compares the Master Files.
- Compare Access Files* compares the Access Files.
- Merge Files* merges the Master Files.
- Merge Access Files* merges the Access Files.

The files open in side-by-side panes. The differences are highlighted in light blue, with arrows indicating insertions and deletions. For Merge Files, the arrows let you copy lines between the files.




Editing a Stored Procedure




When the files first open, left and right scrolling is locked to the first difference, aligned in the middle of the window (when the file sizes allow), as shown in the following image for Merge Files.



If you click an arrow to copy a line from one file to the other, the highlighting and arrow disappear once the line is copied.

You can use the following buttons to move between differences and control the merge process. The arrows for the current difference will be aligned and displayed with left and right scrolling locked.

Button Image	Description
	<p>Save File (Merge only)</p> <p>Each file has its own Save icon (Save Left and Save Right), so that you can save the file when you are finished with the merge edits for each.</p>
	<p>Save File As (Merge only)</p> <p>Each file has its own Save As icon (Save Left As and Save Right As), so that you can save the file in a new folder and/or under a new name when you are finished with the merge edits for each.</p>
	<p>Jump to Previous Difference</p> <p>Locks and aligns the left and right scroll bars to the previous difference, if there is one.</p>

Button Image	Description
	Jump to Next Difference Locks and aligns the left and right scroll bars to the next difference, if there is one.
	Undo (Merge Only) Undoes the previous merge action.
	Redo (Merge Only) Redoes the previously undone merge action.

Running a Stored Procedure

You can run a procedure from the Reporting Server browser interface or the Data Migrator desktop interface with the following variations:

- Run.** Runs the procedure immediately.
- Run Changing Defaults.** Allows you to change variable values prior to running the procedure.
- Run Advanced.** Has the following options:
 - Submit with E-Mail.** Allows you to set up email notification when the procedure starts and/or completes, and to change default values of amper variables.
 - Run Stress.** Records the exact sequence of user actions into a script. For information, see [Stress Testing a Procedure](#) on page 565.
 - Run Stress Deferred.** Records the exact sequence of user actions into a script, but delays the beginning of the run. For information, see [Stress Testing a Procedure](#) on page 565.
 - Debug.** Enables you to run the procedure in debug mode. You can step through the code, viewing the result for each line.

You can run a stored procedure from an application folder in the resources tree of the Applications page.

In addition, when you create or edit a stored procedure in the text editor, you can run the procedure directly from that window, with or without first saving it.

Procedure: How to Run a Procedure

You can run procedures from the Reporting Server browser interface or the Data Migrator desktop interface or from the command line.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.

On the Reporting Server browser interface, the Applications page opens.

2. Expand an application folder.
3. Right-click a procedure, and select *Run*.

The output is displayed in the right pane.

If output cannot be generated, an error message is displayed.

To run a stored procedure from the command line, use the `edastart -x` command. The syntax is

```
edastart -x "EX flex_name"
```

where:

flex_name

Is the name of the stored procedure you want to run.

Note: On Windows, the default directory from which you can run a procedure is:

```
C:\ibi\srvxx\wfs\bin
```

where:

xx

Is the server release number.

Running a Procedure With Different Variables

You can run a stored procedure with amper variable values that differ from those originally specified in the procedure.

How to Run a Procedure With Different Variable Values

You can run procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.

2. Expand an application folder.
3. Right-click a procedure and select *Run changing defaults*.
4. Enter values for the variables and click *Run*.

Please enter value(s) for the following variable(s)

REGION

Note: You can also change variable values by selecting *Run Advanced* and then *Submit with E-Mail*. There is a *Submit changing defaults* check box on the Submit page that opens.

The output is displayed in the right pane.

Executing a Repository FOCEXEC From EDASPROF

Repository procedures can be executed, and Master and Access Files stored in a repository can be accessed, during server, group, and user profile execution after all DBMS connections and an APP MAP command to the repository have been executed.

For information about creating and managing an SQL Repository, see [Managing Applications](#) on page 117.

Prompting for Amper Variables in Procedures

First-level (non-nested) procedures can prompt for Dialogue Manager amper variables with the following exceptions: -IF THEN ELSE branching, -GOTO branching, and -PROMPT with a list of amper variable values. (These commands may produce unexpected results.)

Debugging a Stored Procedure

The Reporting Server browser interface has a Dialogue Manager Debugger tool. Using this tool, you can execute a procedure that contains Dialogue Manager commands while watching and changing the values of the Dialogue Manager variables in the procedure. You can set and remove breakpoints in the procedure, step into each line of code, step over execution of external procedures, continue to a breakpoint, execute an ad-hoc procedure, and open HOLD files created in the procedure. You can save the session information of one debug session.

While the procedure is executing, the Dialogue Manager Debugger shows the values of all local variables, global variables, and, if requested, system variables. If a TABLE request is executed, the debugger shows the FOCUS stack with the commands being executed. If a variable that has not been defined is used in the procedure, the Dialogue Manager Debugger prompts you for a value.


Note that when the Dialogue Manager Debugger is executing a procedure, it is not running that procedure in an interactive WebFOCUS environment, so interactive Dialogue Manager commands (such as -PROMPT) that require a response from the user are not executed. While in a debugging session, you can open multiple procedures and Master Files. Each opens in its own tab.

Using the Dialogue Manager Debugger

The Dialogue Manager Debugger is for use with WebFOCUS procedures that contain Dialogue Manager commands.

To open a procedure in the Dialogue Manager Debugger, right-click the procedure, point to *Run Advanced*, and click *Debug*.

The procedure opens in the Dialogue Manager Debugger, as shown in the following image.



The left panel displays the procedure with line numbers. The Debug panel displays the variables and stacks used in the procedure. Once you start execution, an output panel opens on the right showing all output.

The Debug panel displays the following items.

- Watch Local Amper Variables (&).** Displays values of the local variables you add to the list. Adding variables to this list enables you to change their values at breakpoints.
- All Local Amper Variables (&).** Displays the values of all local variables in the procedure. You cannot change variable values from this list.

Note: If you want to include system variables, once the procedure is running, right-click *All Local Amper Variables (&)*, and click *Show system ampers*.
- Watch Global Amper Variables (&&).** Displays values of the global variables you add to the list. Adding variables to this list enables you to change their values at breakpoints.
- All Global Amper Variables (&&).** Displays the values of all global variables in the procedure. You cannot change variable values from this list.
- Breakpoints.** Lists the breakpoints you set.
- Focus stack.** Lists the WebFOCUS commands being executed, for example, a TABLE request.
- Execution stack.** Lists the procedures being executed. You can right-click a procedure and click *View focexec* to open the FOCEXEC in a tab in the left panel.

The Output panel shows report output and messages. It has an *Output actions* menu with the option *Clear output*, to clear the panel content.

Setting Breakpoints

To set a breakpoint, right-click a Dialogue Manager command line where you want the breakpoint to occur, and click *Set Breakpoint*.

A breakpoint is a stopping point for execution. You may want to set one or more breakpoints in order to examine or change variable values at those points, run an ad-hoc request, or open a HOLD file created in the procedure.

Once you set a breakpoint, a solid circle appears in the Brk column on that line of the procedure, and the breakpoint is added to the list of breakpoints in the Debug panel.

You can disable or clear a breakpoint by right-clicking the line where the breakpoint is set and clicking *Disable Breakpoint* or *Clear Breakpoint*. A disabled breakpoint appears as a white circle with only an outline.

You can also click the *Breakpoint Actions* menu in the Breakpoints section of the Debug panel and click one of the following options from the context menu.

- Add new breakpoint.
- Enable all breakpoints.
- Disable all breakpoints.
- Remove all breakpoints.

If you click *Add new breakpoint*, a dialog box opens for you to enter a procedure name and line number for the breakpoint.

If the breakpoint is for the current procedure, leave the Focexec text box blank.

If you want to set a breakpoint at the entry point of the procedure (whether or not it has a Dialogue Manager command), leave the line number blank.

Click *Submit* to set the breakpoint.

You can also set a breakpoint at the end of a file, even if there is no Dialogue Manager command on the line. Click the *Breakpoint Actions* menu in the Breakpoints section of the Debug panel and click *Add new breakpoint*. Enter the line number, and click *Submit*.

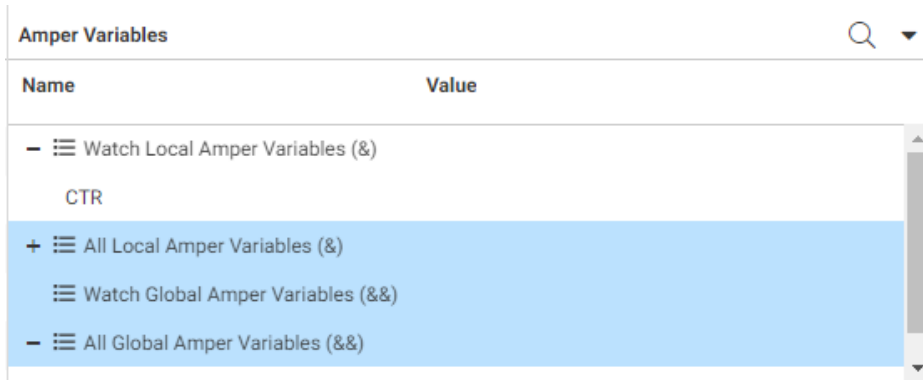
Watching Variables

To add variables to a watch list, right-click *Watch Local Amper Variables (&)* or *Watch Global Amper Variables (&&)*, and click either *Add new amper variable* or *Add all amper variables*.

By default, as variables are set or changed, they are listed under *All Local Amper Variables (&)* or *All Global Amper Variables (&&)*. You can add specific or all variables to the *Watch Local Amper Variables (&)* list or the *Watch Global Amper Variables (&&)* list, if you want to change their values at a breakpoint.

If you click *Add new amper variable*, a dialog box opens in which you enter the name of the variable to watch (without the preceding ampersands).

The variable is added to the watch list, as shown in the following image.



Debugging Steps

To start the debugger, click *Start debugging* on the ribbon.

The commands for executing the procedure and running an ad-hoc request become active.

If the procedure references a variable that has not been given a value, a dialog box opens to prompt you for a value when the variable is encountered.

The following options are available from the ribbon.

File Options

- Open Focexec.** Opens a file picker dialog box that enables you to open a FOEXEC file (.fex).
- Open Master.** Opens a file picker dialog box that enables you to open a Master File (.mas).
- Open Hold File.** Opens a dialog box in which you can enter the name of a HOLD file created in the debugging session as long as the procedure is still running.
- Save Debug Session.** Saves the current debugging session. Only one session can be saved. A new session will overwrite the existing saved session.
- Restore Debug Session.** Restores the saved debugging session.
- Clear Debug Session.** Clears the saved debugging session.
- Switch Viewing File.** If you have opened or executed multiple files, allows you to select among them.
- To close a file, click the *X* on its tab or open the Options menu on the tab and click *Close*. If only one file is open, it cannot be closed without closing the debugger.

- ❑ To close the debugger, along with any files open in the debugger, click the X at the top right of the debugger page.

Execution Options

- ❑ **Start debugging.** Starts the debugging session by activating the other execution buttons.
- ❑ **Step into command.** Steps into commands in an external procedure (invoked by the EX or -INCLUDE commands). Otherwise, Step into executes to the next Dialogue Manager command in the procedure.
- ❑ **Step to next command.** Executes to the next Dialogue Manager command in the procedure, bypassing commands in an external procedure (invoked by the EX or -INCLUDE command).
- ❑ **Continue execution until next breakpoint.** Executes all commands to the next breakpoint.
- ❑ **Execute remainder of current focexec without stopping at breakpoints.** Executes the remainder of the current procedure without stopping at breakpoints.
- ❑ **Stop focexec execution.** Ends the debugging session and stops the execution of all procedures in the execution stack. You can click *Start* to start a new debugging session.
- ❑ **Enter and run ad hoc focexec.** Opens a window for entering and running a FOCEXEC.
Enter a value and click *Submit*.

The following options are available on the Options menu for a file open in a tab:

- ❑ **Scroll to line number.** Opens a dialog box in which you can enter a line number. When you click *OK*, the file is scrolled to the specified line number.
- ❑ **Edit.** Opens the FOCEXEC or Master File in the text editor.
- ❑ **Close.** Closes the file.

Editing a FOCEXEC or Master File in the Debugger

On the *Focexec/Master actions* menu of a FOCEXEC or Master File tab, click *Edit*.

The file opens in the text editor, as shown in the following image.

```

baseapp/dmdebugger3.fex
1 |SET &CTR = 1;
2 -TYPE REQUEST # &CTR;
3 *TABLE FILE WF_RETAIL_LITE
4 SUM DAYSDELAYED AS DAYS
5 BY TIME_MTH
6 WHERE BUSINESS_REGION EQ '&REGION'
7 WHERE_GROUPED DAYSDELAYED GT 200
8 ON TABLE SET PAGE NOPAGE
9 END
10 -TYPE END OF REQUEST &CTR;
11 -RUN
12 -SET &CTR = &CTR+1;
13 -TYPE REQUEST # &CTR;
14 -INCLUDE doc77x/catfex2
15 -TYPE END OF REQUEST &CTR;
16 -RUN
17 -SET &CTR = &CTR+1;
18 -TYPE REQUEST # &CTR;
19 EX baseapp/catfex3
20 -TYPE END OF REQUEST &CTR;
21 -RUN
22 -TYPE END OF PROCEDURE
23
24
25
26

```

For information about using the text editor, see [Editing a Stored Procedure](#) on page 459.

Once you have made your edits, click **Save**. A message similar to the one in the following image displays.

Warning



Focexec baseapp/dmdebugger3 will be out of sync with debugger. Debug session will be restarted.
Continue?

When you click *OK*, the FOCEXEC or Master File will be saved, and the debugging session will be restarted.

Note: Depending on the edits you make to a procedure, you may have to adjust your breakpoints so that they occur on the lines at which you want to break.

Dialogue Manager Debugger Sample Session

The following procedure is named `dmdbugger3.fex`. It includes a variable named `&CTR` whose value is set in the procedure, and a variable named `®ION` whose value is not set in the procedure. The procedure contains a `TABLE` request and two external requests. The first external request displays output and the second creates a `HOLD` file named `cathold`.

```
-SET &CTR = 1;
-TYPE REQUEST # &CTR;
TABLE FILE WF_RETAIL_LITE
SUM  DAYSDELAYED AS DAYS
BY  TIME_MTH
WHERE BUSINESS_REGION EQ '&REGION'
WHERE_GROUPED DAYSDELAYED GT 200
ON TABLE SET PAGE NOPAGE
END
-RUN
-TYPE END OF REQUEST &CTR;
-SET &CTR = &CTR+1;
-TYPE REQUEST # &CTR;
EX baseapp/catfex2
-RUN
-TYPE END OF REQUEST &CTR;
-SET &CTR = &CTR+1;
-TYPE REQUEST # &CTR;
EX baseapp/catfex3
-RUN
-TYPE END OF REQUEST &CTR;
-TYPE END OF PROCEDURE
```

The following image shows this procedure open in the debugger, with breakpoints set on Dialogue Manager commands before and after each request. The breakpoints are listed on the Debug panel.

The screenshot displays the debugger interface for a stored procedure. The main window is titled 'baseapp/dmdebugger3' and contains a table of code lines. The right-hand side of the interface is divided into several panels: 'Amper Variables', 'Breakpoints', 'Focus stack', and 'Execution Stack'.

Code Lines:

Brk	Line	Focexec
	1	-SET &CTR = 1;
●	2	-TYPE REQUEST # &CTR;
	3	TABLE FILE WF_RETAIL_LITE
	4	SUM DAYSDELAYED AS DAYS
	5	BY TIME_MTH
	6	WHERE BUSINESS_REGION EQ 'REGION'
	7	WHERE_GROUPED DAYSDELAYED GT 200
	8	ON TABLE SET PAGE NOPAGE
	9	END
	10	-TYPE END OF REQUEST &CTR;
●	11	-RUN
	12	-SET &CTR = &CTR+1;
●	13	-TYPE REQUEST # &CTR;
	14	-INCLUDE doc77x/catfex2
●	15	-TYPE END OF REQUEST &CTR;
	16	-RUN
	17	-SET &CTR = &CTR+1;
	18	-TYPE REQUEST # &CTR;
	19	EX baseapp/catfex3
●	20	-TYPE END OF REQUEST &CTR;
	21	
	22	-RUN
	23	-TYPE END OF PROCEDURE
	24	
	25	
	26	

Amper Variables Panel:

Name	Value
Watch Local Amper Variables (&)	
All Local Amper Variables (&)	
Watch Global Amper Variables (&&)	
All Global Amper Variables (&&)	

Breakpoints Panel:

Focexec Name	Line number	Enabled
baseapp/dmdebugger3	2	yes
baseapp/dmdebugger3	11	yes
baseapp/dmdebugger3	13	yes
baseapp/dmdebugger3	15	yes
baseapp/dmdebugger3	20	yes

Execution Stack Panel:

Focexec Name	Line number
--------------	-------------

Clicking *Start*, then *Step to next command* stops at the first breakpoint. Click the plus signs (+) next to *All Local Amper Variables (&)* and *All Global Amper Variables (&&)*. The value of the &CTR local variable is shown as 1. The value of the &&GEO_UNIFIED_ROLE global variable is shown as ON. The execution stack shows the name and line number of the procedure being executed, as shown in the following image.

The screenshot displays a debugger window for 'baseapp/dmdebugger3'. The main pane shows a list of code lines with breakpoints. Line 2 is highlighted, and a search icon is visible in the top right of this pane. The right-hand pane is divided into several sections:

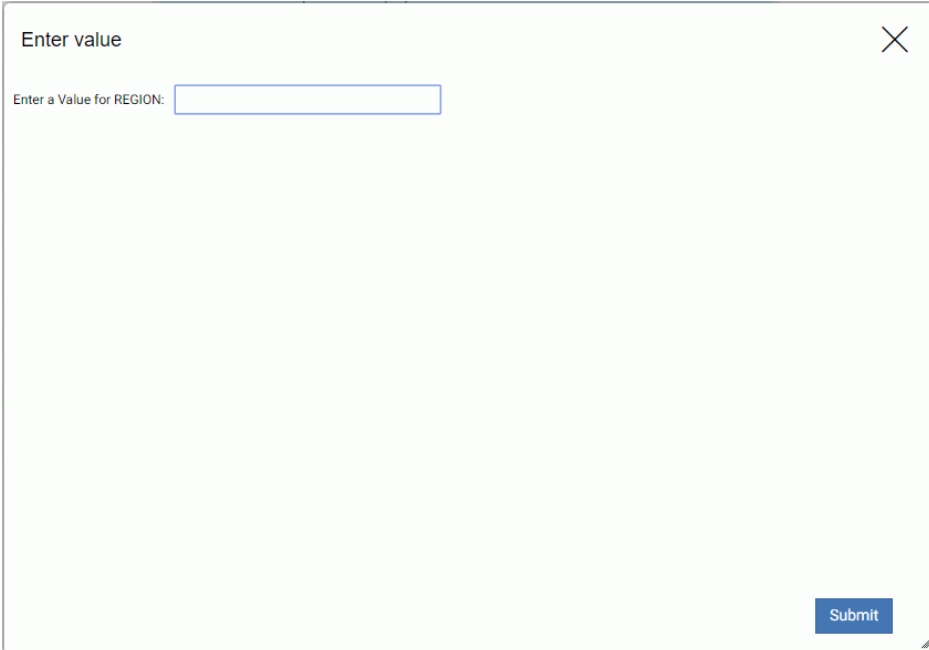
- Amper Variables:** A table showing local and global variables.

Name	Value
+ Watch Local Amper Variables (&)	
- All Local Amper Variables (&)	
CTR	1
+ Watch Global Amper Variables (&&)	
- All Global Amper Variables (&&)	
GEO_UNIFIED_ROLE	ON
- Breakpoints:** A table listing active breakpoints.

Focexec Name	Line number	Enabled
baseapp/dmdebugger3	2	yes
baseapp/dmdebugger3	11	yes
baseapp/dmdebugger3	13	yes
baseapp/dmdebugger3	15	yes
baseapp/dmdebugger3	20	yes
- Focus stack:** Shows 'EMPTY'.
- Execution Stack:** Shows the current execution context.

Focexec Name	Line number
#0 baseapp/dmdebugger3	2

Clicking *Step into command* opens the following dialog box that asks you to enter a value for the ®ION variable.



Enter value

Enter a Value for REGION:

Submit

Entering EMEA and clicking *Submit* proceeds to the next Dialogue Manager command. REGION is added to the list of local variables, and the first message displays in the output window. Also, before the request is actually run, the TABLE commands to be executed are listed on the Focus Stack, as shown in the following image.

The screenshot shows a debugger window titled 'baseapp/dmdebugger3'. The interface is split into three main sections:

- Left Panel (Focexec List):** A list of commands with columns 'Brk', 'Line', and 'Focexec'. Line 10 is highlighted with a blue background.


Brk	Line	Focexec
	1	-SET &CTR = 1;
●	2	-TYPE REQUEST # &CTR;
	3	TABLE FILE WF_RETAIL_LITE
	4	SUM DAYSDELAYED AS DAYS
	5	BY TIME_MTH
	6	WHERE BUSINESS_REGION EQ 'REGION'
	7	WHERE_GROUPED DAYSDELAYED GT 200
	8	ON TABLE SET PAGE NOPAGE
	9	END
⊕	10	-TYPE END OF REQUEST &CTR;
●	11	-RUN
	12	-SET &CTR = &CTR+1;
●	13	-TYPE REQUEST # &CTR;
	14	-INCLUDE doc77x/cattex2
●	15	-TYPE END OF REQUEST &CTR;
	16	-RUN
	17	SET &CTR = &CTR+1;
	18	-TYPE REQUEST # &CTR;
	19	EX baseapp/cattex3
●	20	-TYPE END OF REQUEST &CTR;
	21	
	22	-RUN
	23	-TYPE END OF PROCEDURE
	24	
	25	
	26	
- Right Panel (Amper Variables):** A table showing variable names and their values.

Name	Value
- Watch Local Amper Variables (&)	
CTR	1
- All Local Amper Variables (&)	
CTR	1
REGION	EMEA
- Watch Global Amper Variables (&&)	
- Bottom Panel (Focus Stack):** A list of commands currently in focus.

Focexec Name	Line number	Enabled
baseapp/dmdebugger3	2	yes
baseapp/dmdebugger3	11	yes
baseapp/dmdebugger3	13	yes
baseapp/dmdebugger3	15	yes
baseapp/dmdebugger3	20	yes

After the request is run, the output displays in the output window. The next request in the procedure is issued using a `-INCLUDE` command. The external procedure is opened in the debugger if the last debugger command was `Step into` command, or the last debugger command was `Step to next` command and the external procedure contains a breakpoint on a line that will be executed.

The last request in the procedure, run using an `EX` command, creates a `HOLD` file named `cathold`. After the request is run, but while the procedure is still running, you can open the `HOLD` file by clicking *Open Hold File* from the File menu, as shown in the following image.



A dialog box opens where you can enter the name of the hold file.

After you enter the name and click *Submit*, a report showing the HOLD file data displays in the output window, as shown in the following image.



BRAND REVENUE_US
Audio Technica \$51,694.24
BOSE \$221,651.35
Canon \$146,110.76
Denon \$35,419.48
Grado \$27,695.75
Harman Kardon \$98,300.75
JVC \$263,957.51
LG \$235,056.84
Logitech \$81,626.78
Niles Audio \$103,915.15
Onkyo \$169,921.05
Panasonic \$398,330.00
Philips \$66,707.67
Pioneer \$312,858.54
Polk Audio \$83,271.27
Roku \$17,358.73
Samsung \$295,142.54
Sanyo \$94,548.36
Sennheiser \$112,887.86
Sharp \$127,069.77

To clear the output pane, click the *Output actions* menu, and click *Clear output*.

After the procedure finishes executing, the debugger stops. You can click *Start* to start a new execution.

Running a Procedure With Advanced Options

You can submit a procedure and specify email options to send an email at start, at completion, and at failure. You can also run a stress test.

To run a procedure with advanced options:

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, select *Run Advanced*, and then either *Submit with E-Mail*, *Run Stress*, or *Run Stress Deferred*.

If you selected *Submit with E-Mail*, the following page opens.

The screenshot shows a 'Submission' dialog box with the following sections and options:

- Submission** (expanded):
 - Run at a specified date and time
 - Do not wait for completion
 - Submit changing defaults
 - Submit with Limits
 - Report name:
- Log and output destinations** (expanded):
 - E-Mail On Start** (expanded)
 - E-Mail On Completion** (expanded):
 - Mail to: Defines where to send E-mail
 - To:
 - Cc:
 - Bcc:
 - Use a semicolon as the separator character for multiple email addresses (id@domain.com), user IDs (myid) or procedure names (%myfex).
 - Advanced Advanced E-Mail Settings on Completion
 - E-Mail On Failure** (expanded)
- Buttons:

Check the submission and log options you want, fill in the email fields if you want to send emails at start or completion, and click *Submit*.

- ❑ If you selected *Run Stress* or *Run Stress Deferred*, the following page opens.

The screenshot shows a configuration window with a 'Basic' tab. The fields are as follows:

Script Directory:	doc77x
Script Name:	dmdebugger3
? Number Of Threads:	1
? Interval Parameter:	0.5
? Keep alive for:	<input type="text"/> seconds
? Delay:	<input type="text"/> seconds between requests in keep-alive mode
? Test Description:	<input type="text"/>

Below the 'Basic' section is an 'Advanced' section which is collapsed. At the bottom of the window is a 'Start' button.

You can fill in the fields to create a test, and click *Start* to run it, or use the buttons at the top of the window to create a new script file or recording.

For information on running a stress test, see [Stress Testing a Procedure](#) on page 565.

Scheduling a Procedure

You can control when and how often a procedure is run using the *Schedule and E-mail* option on the resources tree of the Applications page.

Procedure: How to Schedule a Request to Run Once

You can schedule procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, select *Schedule and E-mail*, and then *Manage*.

The Manage Schedule and E-mail for procedure page 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 Send E-Mail on Completion
 ? Mail to To Defines where to send E-mail
 Cc
 Bcc
 Use a semicolon as the separator character for multiple email addresses (id@domain.com), user IDs (myid) or procedure names (%myfex).
 ? Advanced Advanced E-Mail Settings on Completion
 ^ E-Mail On Failure
 Cancel Set Set and Test Run
 Create new server schedule object to submit this procedure Save

4. Select *Active* from the Schedule Status drop-down menu.

The Start and stop range section opens, as shown in the following image.

^ General
 ? Schedule Status Active Enables/Disables Schedule
 ? Schedule Type Run Once Determines whether a flow will execute once, on a recurring basis, or on several specific days
 ^ Start and stop range
 ? Start Date 03/27/2018 Specifies the start time and date for the execution of the flow
 ? Start Time 10:03 Specifies the start time and date for the execution of the flow
 ^ Special date/time ranges
 ? Specify Stop Time Specifies forced end time

5. To set a start date, click the *Start Date* drop-down menu to open a calendar, and select a date.
6. To select a start time, select an hour and minute from the Start Time drop-down menus. Hours are listed for a 24-hour clock.
 Optionally, you can set a forced stop time.
7. Optionally, you can set up email notification when the procedure starts and/or ends in the E-Mail on Start and E-Mail on Completion sections.

Note: To use email notification, the E-mail SMTP Server must be configured from the Workspace page. For information, see [SMTP Server Settings](#) on page 617.

You can also specify options for logging.

8. Click *Set* or *Set and Test Run*.

Procedure: How to Schedule a Recurring Request to Run

You can schedule procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, select *Schedule and E-mail*, and then *Manage*.

The Manage Schedule and E-mail for procedure page opens.

4. Select *Active* from the Schedule Status drop-down menu.
5. In the Schedule Type drop-down list, select *Recurring*.

Additional parameters become available in the General Section, as shown in the following image.

^ General		
? Schedule Status	Active ▾	Enables/Disables Schedule
? Schedule Type	Recurring ▾	Determines whether a flow will execute once, on a recurring basis, or on several specific days
? Interval Type	Hourly ▾	Specifies the interval in which a flow will be executed
? Interval Value	1	Specifies the frequency at which a flow will be executed

6. Select an interval from the Interval Type drop-down menu. The choices are *Minutes*, *Hourly*, *Daily*, *Weekly*, *Monthly* and *Yearly*.
7. Enter a value in the Interval value field. This specifies the frequency at which the procedure will run. For example, if you chose *Weekly* as the Interval Type, entering 2 will execute the procedure every two weeks.
8. To set a start date, click the *Start Date* drop-down menu to open a calendar, and select a date.
9. To select a start time, select an hour and minute from the Start Time drop-down menus. Hours are listed for a 24-hour clock.
10. Optionally, you can select the *Specify End Date* check box, and select an end date and time.

11. Optionally, you can select days of the week or month, intraday start and/or end, as well as how to handle holidays, from the Special date/time ranges section, as shown in the following image.

Special date/time ranges								
? Days of the Week	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Specifies day or days (with multi-select) of week to run this flow
? Days of the Month	1	2	3	4	5	6	7	Specifies day or days (with multi-select) of month to run this flow
	8	9	10	11	12	13	14	
	15	16	17	18	19	20	21	
	22	23	24	25	26	27	28	
	29	30	31	Last Day of Month				
? Intraday Start	00	00						Specifies the beginning of the daily time span
? Intraday End	00	00						Specifies the ending of the daily time span
? Run on Holidays							Specifies option to run on holidays	

The Run on Holidays drop-down menu allows you to skip, run, or only run the report on holidays.

12. Optionally you can expand the Log and output destinations section and select destinations, as shown in the following image.

Log and output destinations	
? ETLLLOG	<input checked="" type="checkbox"/> Send log to ETLLLOG and output to disk
? EDAPRINT	<input type="checkbox"/> Send log to EDAPRINT
? E-Mail	<input checked="" type="checkbox"/> Send log and output as attachments to E-mail if E-mail Notification 'On Completion' or 'On Failure' is requested
? Send HTML output as inline message	<input type="checkbox"/> Embed HTML output into E-Mail message

13. Optionally, you can set up email notification when the procedure starts and/or ends in the E-Mail on Start and E-Mail on Completion sections.

Note: To use email notification, the E-mail SMTP Server must be configured from workspace.

14. Click *Set* or *Set and Test Run*.

Procedure: How to Schedule a Request to Run on Multiple Days

You can schedule procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, point to *Schedule and E-mail*, and click *Manage*.

The Manage Schedule and E-mail for procedure page opens.

4. Select the *Active* from the Scheduler Status drop-down menu.
5. In the General section, select *Multi-Day* from the Schedule Type drop-down menu.
The Special date/time ranges section opens, as shown in the following image.

Special date/time ranges

? Days of the Week: Sun Mon Tue Wed Thu Fri Sat
Specifies day or days (with multi-select) of week to run this flow

? Days of the Month: 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 Last Day of Month
Specifies day or days (with multi-select) of month to run this flow

? Intraday Start: 00 00
Specifies the beginning of the daily time span

? Intraday End: 00 00
Specifies the ending of the daily time span

? Run on Holidays: [dropdown]
Specifies option to run on holidays

6. To set a start date, click the *Start Date* drop-down menu to open a calendar, and select a date.
7. To select a start time, select an hour and minute from the Start Time drop-down menus. Hours are listed for a 24-hour clock.
8. Optionally, you can select the *Specify End Date* check box, and select an end date and time.
9. Optionally, you can select days of the week or month, as well as how to handle holidays, from the Special date/time ranges section.

The Run on Holidays drop-down menu allows you to skip, run, or only run the report on holidays.

10. Optionally you can expand from the Log and output destinations section and select destinations, as shown in the following image.

Log and output destinations

? ETLLOG Send log to ETLLOG and output to disk

? EDAPRINT Send log to EDAPRINT

? E-Mail Send log and output as attachments to E-mail if E-mail Notification 'On Completion' or 'On Failure' is requested

? Send HTML output as inline message Embed HTML output into E-Mail message

11. Optionally, you can set up email notification when the procedure starts and/or ends in the E-Mail on Start and E-Mail on Completion sections.

Note: To use email notification, the E-mail SMTP Server must be configured from the Workspace page.

12. Click *Set* or *Set and Test Run*.

Procedure: How to Schedule a Request to Run When the Server Starts

You can schedule procedures from the Reporting Server browser interface or the Data Migrator desktop interface. This procedure will run on every server start.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, point to *Schedule and E-mail*, and click *Manage*.

The Manage Schedule and E-mail for procedure page opens.

4. Select the *Active* from the Schedule Status drop-down menu.
5. In the General section, select *Run when server starts* from the Schedule Type drop-down menu.
6. Optionally, you can set up an email notification when the procedure starts and/or ends in the E-Mail on Start and E-Mail on Completion sections.

Note: To use the email notification option, the E-mail SMTP Server must be configured from the workspace.

7. Optionally, you can expand from the Log and output destinations section and select destinations.
8. Click *Set* or *Set and Test Run*.

Sending an E-mail Notification for a Procedure

You can set up your server to send an email message that notifies recipients of the successful completion or failure of every procedure execution. In addition, you can request that an email message be sent to a designated recipient at the start or end of the request execution.

Prerequisites: This procedure assumes that the E-mail SMTP Server has been configured. For information, see [SMTP Server Settings](#) on page 617.

Procedure: How to Send an E-mail Message for a Request

You can set up email for procedures from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.

On the Reporting Server browser interface, the Applications page opens.

2. Expand an application folder.

3. Right-click a procedure, point to *Schedule and E-mail*, and click *Manage*.

The Manage Schedule and E-mail for procedure page opens.

4. Optionally, expand the E-Mail on Start section and enter one or more email addresses in the *Mail to* field. You can optionally send carbon copies and blind carbon copies using the Cc and Bcc fields.
5. In the E-Mail on Completion section, enter an email address in the *Mail to* field. You can optionally send carbon copies and blind carbon copies using the Cc and Bcc fields.
6. Optionally, select the *Advanced* check box in either E-Mail section.

Additional parameters become available, as shown in the following image.

E-Mail On Completion

? Enable

? Mail to

Cc

Bcc

Use a semicolon as the separator character for multiple email addresses (id@domain.com).

? Advanced

? Mail to addresses defined in procedure

? Importance

? Subject

? E-Mail Message

Send E-Mail on Completion
Defines where to send E-mail

Advanced E-Mail Settings on Completion
Specifies procedure defining where to send E-mail
Specifies the importance of the E-Mail message
Specifies the subject of the E-Mail message
Enter message to be included into E-Mail body

7. To provide advanced parameters:

- Mail to addresses defined in procedure* allows you to select a procedure that contains the list of email addresses.
- Select an option from the Importance drop-down menu, either *Normal*, *Low* or *High*.
- Enter information in the *Subject* field.
- Enter in information in the *E-Mail Message* field.
- Select the *Send HTML output as inline message* check box in the Log and Output Destinations section (only available for E-Mail on Completion) to use the report as the email body.

Note: Any content in the *E-Mail Message* field will not be included.

8. Click *Set* or *Set and Test Run*.

Procedure Run-Time Processing Essentials


A stored procedure is a file of commands that resides on a server and is also known as a FOCEXEC procedure. The procedure typically includes any combination of Dialogue Manager control statements, Dialogue Manager variables, SET commands, SQL syntax, TABLE syntax and other syntax that perform tasks such as report generation or file maintenance, or it simply generates messages.

This section is intended as a basic reference for Dialogue Manager and how it integrates to some basic TABLE tasks. Complete information about Dialogue Manager, see the *TIBCO WebFOCUS® Stored Procedure and Subroutine Reference for 3GL Languages* manual and *Managing Flow of Control in an Application in the TIBCO WebFOCUS® Developing Reporting Applications* manual.

Commands in a Stored Procedure

A procedure must reside on disk in an APP location, which may be an APP directory or an APP Mapped directory (or PDS on z/OS PDS deployment). Additionally, a procedure may reside elsewhere and be "sent" to the server by a client application such as the WebFOCUS mid-tier.

With the EXEC (or EX) command, a stored procedure may be called by another stored procedure or by a client application. This is illustrated below.



In addition to Dialogue Manager commands that control execution flow, and messages (such as the -TYPE above), the following may be included in a procedure:

- ❑ SQL and TABLE statements allowed by the server platform.
- ❑ Server commands, for example, CALLPGM, EXEC, and END. For details on CALLPGM and EXEC when used to call a compiled 3GL program, see the *TIBCO WebFOCUS® Stored Procedure and Subroutine Reference for 3GL Languages* manual. This chapter discusses the use of EXEC to call other FOCEXEC procedures.
- ❑ Commands allowed in a server profile, such as SET. For details on the profile and its allowable commands, see *TIBCO WebFOCUS® Reporting Server Introduction* on page 17.
- ❑ Commands that enable portions of a procedure to be executed on a remote alternate server. See the *TIBCO WebFOCUS® Stored Procedure and Subroutine Reference for 3GL Languages* manual for details on the syntax and use of those commands.

- ❑ The ON TABLE HOLD command, which holds an answer set in a temporary file on a server. See [How to Use the ON TABLE HOLD Command](#) on page 518 for details on syntax and use.
- ❑ The ON TABLE PCHOLD command, which sends an answer set to a client application. See [Use the ON TABLE PCHOLD Command](#) on page 519 for details on the syntax and use.
- ❑ Platform-specific commands such as file allocation (that is, DYNAM on z/OS or FILEDEF on other platforms).

Dialogue Manager Commands and Processing

The following table summarizes the available Dialogue Manager commands. Note that every command begins with a hyphen (-), and this is what differentiates Dialogue Manager syntax from other syntax within a procedure.

The following sections describe the syntax and use of the commands. For an alphabetical list, see the *TIBCO WebFOCUS® Stored Procedure and Subroutine Reference for 3GL Languages* manual.

Command	Function
-*[<i>text</i>]	Comments within a procedure.
-?	Displays the value of local variables.
-UNIX <i>cmd</i> -WINNT <i>cmd</i> -DOS <i>cmd</i> -AS/400 <i>cmd</i> -TSO RUN <i>cmd</i> -SYSTEM <i>cmd</i>	Executes an operating system command for the respective operating system, and the command is ignored when encountered on a non-applicable operating system, except -SYSTEM which runs on all platforms. Note that some platforms use a prior version of the platform name (that is, -AS/400 is for IBM i and -WINNT is for Windows). There are also non <i>dash</i> versions of these commands (and ! is the equivalent of -SYSTEM), but they are not ignored when run on a platform for which they are not relevant.
-CLOSE	Closes an external file opened for reading or writing (an external file is a sequential file in the platform's file system).

Command	Function
-DEFAULT -DEFAULTS	Sets a variable to an initial value.
-EXIT	Executes stacked commands and terminates the procedure. For a definition of stacked commands, see Commands in a Stored Procedure on page 495.
-GOTO	Forces an unconditional branch to a label.
-IF	Determines the execution flow based on the evaluation of an expression (a conditional branch).
-INCLUDE	Calls another Dialogue Manager FOCEXEC procedure.
-label	Identifies a section of code that is the target of a -GOTO or -IF.
-PASS	Directly issues and controls passwords.
-PROMPT	Types a message to the terminal (if edastart -t is in use) or creates an input window with the message in a browser if the connection type is HTTP and reads the reply from the user. This reply assigns a value to the variable named.
-QUIT	Terminates the procedure without executing stacked commands.
-READ	Reads data from an external file.
-REMOTE BEGIN	Signals the start of commands on an originating server that are to be sent to a target server. Only available with Hub Services.
-REMOTE END	Signals the end of commands from an originating server.


Command	Function
<code>-REPEAT</code>	Executes a loop.
<code>-RUN</code>	Executes stacked commands and closes any external files opened with <code>-READ</code> or <code>-WRITE</code> .
<code>-SET</code>	Sets a variable to a literal value or to a value computed in an expression.
<code>-TYPE</code>	Sends a message to a client application.
<code>-WRITE</code>	Writes data to an external file.
<code>-</code>	Line continuation of a prior Dialogue Manager command.

Dialogue Manager Command Processing in a Procedure

A procedure processes as follows:

- Dialogue Manager reads each line of the procedure, one by one. Values are substituted for variables encountered in any line.
- All Dialogue Manager commands (commands that start with a "-") execute as they are encountered.
- Other commands are temporarily stored (stacked) for subsequent execution and are referred to as *stacked* commands.

- ❑ The Dialogue Manager commands -RUN and -EXIT force the execution of any stacked commands.



Example: Issuing an EXEC (or EX) Call

The following is an example of executing a procedure (SLRPT), with an explanation of the way it processes.

To execute this procedure, the requesting application issues an EXEC (or EX) with the procedure name and supplies values for the required parameters. For example:

```
EX SLRPT COUNTRY=ENGLAND , CAR=JAGUAR
```

The following procedure has been enumerated into sections, and its behavior explained (the numbers on the left are not part of the actual procedure, but are simply supplied as reference points for the explanations):

```
1. -IF &COUNTRY EQ 'DONE' THEN GOTO GETOUT;
2. SQL
   SELECT COUNTRY, CAR, MODEL, BODY
   FROM CAR
   WHERE COUNTRY= '&COUNTRY' AND CAR= '&CAR'
   ORDER BY CAR;
3. TABLE
   ON TABLE PCHOLD
   END
4. -RUN
5. -EXIT

   -GETOUT
   -TYPE NO PROCESSING DONE: EXITING SP
```

The procedure processes as follows:

1. Values for the variables &COUNTRY and &CAR are passed to the procedure before the first line executes. Dialogue Manager substitutes the value ENGLAND for the variable &COUNTRY in the first line and tests for the value DONE. The test fails, so Dialogue Manager proceeds to the next line.

If the value were DONE instead of ENGLAND, control would pass to the label -GETOUT, and the message NO PROCESSING DONE: EXITING SP would be sent to the client application. (Dialogue Manager would skip the intervening lines of code.)

2. The next five lines are SQL. Dialogue Manager scans each for the presence of variables, substituting the value ENGLAND for &COUNTRY and the value JAGUAR for &CAR (remember, those values were passed by EDARPC). As each line is processed, it is placed on a stack to be executed later by the server.
3. The command ON TABLE PCHOLD sends the answer set to the client application.

The command END delimits ON TABLE PCHOLD.

After Dialogue Manager processes the command END, the stacked commands look like this:

```
SQL
SELECT COUNTRY, CAR, MODEL, BODY
FROM CAR
WHERE COUNTRY= 'ENGLAND' AND CAR= 'JAGUAR'
ORDER BY CAR;
TABLE
ON TABLE PCHOLD
END
```

The next line is then processed by the Dialogue Manager.

4. The Dialogue Manager command -RUN forces the stacked commands to execute.
5. The Dialogue Manager command -EXIT terminates the procedure.

Using Variables

This section describes how to use variables in a procedure.

Variables fall into two categories:

- Local and global variables, whose values must be supplied by the procedure at run time.
- System and statistical variables, whose values are automatically supplied by the system when referenced.

The following features apply to all variables:

- A variable stores numbers or a string of text, and is placed anywhere in a procedure.
- A variable refers to a command, a database field, a verb, or a phrase.
- The maximum number of variables allowed in a procedure is 1,024. Because approximately 30 are reserved for server use, the maximum number of user-named variables allowed in a procedure is 994.

Naming Conventions

This section describes how to use variables in a procedure.

Variables fall into two categories:

- Local and global variables, whose values must be supplied by the procedure at run time.
- System and statistical variables, whose values are automatically supplied by the system when referenced.

Local and global variable names are user-defined, while system and statistical variables have predefined names.

The following rules apply to the naming of local and global variables:

- A local variable name is always preceded by an ampersand (&).
- A global variable name is always preceded by a double ampersand (&&).
- Embedded blanks are not permitted.
- If an anticipated value for a variable might contain an embedded blank, enclose the variable in single quotation marks when you refer to it.

- ❑ A variable name may be any combination of the characters A through Z, 0 through 9, and the underscore (_). The first character of the name must be A through Z.
- ❑ Assign a number to a variable, instead of a name, to create a positional variable.

Syntax: **How to Use Variables in a Procedure**

&[&] name

where:

&name

Is the user-defined name of a local variable. The first character of *name* must be A through Z.

&&name

Is the user-defined name of a global variable. The first character of *name* must be A through Z.

The following variables are properly named:

```
&WHICHPRODUCT
&WHICH_CITY
'&CITY'
&&CITY
```

The following variables are improperly named for the reason given:

Invalid	Reason
<code>&WHICH CITY</code>	Contains embedded blank.
<code>&WHICH-CITY</code>	Contains a hyphen (-).
<code>WHICHCITY</code>	Leading ampersand(s) is missing.

Local Variables

Once supplied, values for local variables remain in effect throughout a single procedure. The values are lost after the procedure finishes processing and are *not* passed to other procedures that contain the same variable name.

Example: Using Local Variables

Consider the following procedure in which &CITY, &CODE1, and &CODE2 are local variables.

```

.
.
.
SQL
SELECT SUM (UNIT_SOLD),
        SUM (RETURNS)
FROM SALES
WHERE CITY = '&CITY'
AND PROD_CODE >= '&CODE1'
AND PROD_CODE <= '&CODE2'
.
.
.

```

Assume you supply the following values when you call the procedure:

```
CITY=STAMFORD, CODE1=B10, CODE2=B20
```

Dialogue Manager substitutes the values for the variables as follows:

```

.
.
.
SQL
SELECT SUM (UNIT_SOLD),
        SUM (RETURNS), CITY
FROM SALES
WHERE CITY = STAMFORD
AND PROD_CODE >= B10
AND PROD_CODE <= B20
GROUP BY CITY, PROD_CODE
.
.
.

```

After the procedure executes and terminates, the values STAMFORD, B10, and B20 are lost.

Global Variables

Once a value is supplied for a global variable, it remains in effect throughout the session of a processing service, unless cleared by the server. All procedures that contain the same global variable name receive the supplied value until you terminate the session.

Example: Using Global Variables

The following example illustrates the use of three global variables: &&CITY, &&CODE1, &&CODE2.

```

.
.
.
SQL
SELECT SUM (UNIT_SOLD) ,
        SUM (RETURNS)
FROM SALES
WHERE CITY = &&CITY
AND PROD_CODE >= &&CODE1
AND PROD_CODE <= &&CODE2
.
.
.

```

Predefined System Variables

There are predefined run-time system variables that you can use in a procedure. Dialogue Manager automatically supplies values for these variables whenever the variables are encountered.

The list within this section shows only the commonly used variables. Additional variables exist for other specific features and are documented in the sections for those features.

The Reporting Server browser interface Workspace page provides a complete list of system variables and global variables. For information, see [Core Engine System Variables](#) on page 608 and [Core Engine Global Variables](#) on page 609.

Unless otherwise noted in the table, you may override run time values by replacing or adding the parameter values specified.

System Variable	Description	Format or Value
&APPROOT	Physical location of the APPROOT directory.	Directory name.
&DATE	Current date.	MM/DD/YY
&DATE <i>fmt</i>	Current date.	fmt is any combination of YYMD, MDYY, and so on.
&MDY	Current date. Useful for numerical comparisons.	MMDDYY

System Variable	Description	Format or Value
<code>&MDYY</code>	Current date (four-digit year).	<code>MMDDCCYY</code>
<code>&DMY</code>	Current date.	<code>DDMMYY</code>
<code>&DMYY</code>	Current date (four-digit year).	<code>DDMMCCYY</code>
<code>&YMD</code>	Current date.	<code>YYMMDD</code>
<code>&YYMD</code>	Current date (four-digit year).	<code>CCYYMMDD</code>
<code>&FOCFOCEXEC</code>	Current running procedure.	Manages reporting operations involving many similarly named requests that are executed using EX. <code>&FOCFOCEXEC</code> enables you to easily determine which procedure is running. <code>&FOCFOCEXEC</code> is specified within a request or in a Dialogue Manager command to display the name of the currently running procedure.
<code>&FOCINCLUDE</code>	Current included procedure.	Manages reporting operations involving many similarly named requests that are included using <code>-INCLUDE</code> . <code>&FOCINCLUDE</code> is specified within a request or in a Dialogue Manager command to display the name of the current included procedure.
<code>&ECHO</code>	Current echo tracing value.	<code>ON</code> , <code>OFF</code> , or <code>ALL</code>
<code>&FOCMODE</code>	Operating environment.	<code>AS/400</code> , <code>MSO</code> , <code>UNIX</code> , or <code>WINNT</code> . You cannot override the system-supplied value. <code>MSO</code> is the value returned on z/OS. If it is necessary to further test for server PDS versus HFS Deployment. The following may be used: <pre>-? SET TEMP &XYZ -SET &XYZ = IF &XYZ EQ 'MVS' THEN 'PDS' ELSE 'HFS' ;</pre>
<code>&FOCPRINT</code>	Current print setting.	<code>ONLINE</code> <code>OFFLINE</code>

System Variable	Description	Format or Value
&FOCREL	Source code release number.	Release number.
&IORETURN	Value returned after the last Dialogue Manager -READ or -WRITE operation.	<p>0 Successful operation.</p> <p>1 End or failure.</p>
&RETCODE	<p>Return code value from execution of a server or operating system command.</p> <p>Referencing &RETCODE forces the execution of all stacked commands, like the -RUN command.</p>	<p>Any value returned by a command is valid (for example, CALLPGM flag values), but zero is considered normal (successful) execution.</p> <p>The one exception is the &RETCODE value of dash operating system commands, such as -DOS, -UNIX, -AS/400, and -WINNT, represent the success, not of the command they are running, but of the ability of the server to spawn out to the OS and run the command. In this case, the &RETCODE value is normally zero because it reflects that the spawn executes normally regardless of the results of the specific command. For this case, the amper variable &EXITRC should be used to check the command result or the non-dash version of the command should be used.</p>
&EXITRC	<p>Return code value from execution of an operating system command.</p> <p>Referencing &EXITRC forces the execution of all stacked commands, like the -RUN command.</p>	Any value returned by a command is valid, but zero is considered normal (successful) execution.

System Variable	Description	Format or Value
&TOD	Current time. When you enter FOCUS, this variable is updated to the current system time only when you execute a MODIFY, SCAN, or FSCAN command. To obtain the exact time during any process, use the HHMMSS subroutine.	HH.MM.SS
&FOCCODEPAGE	Code page of the server when NLS is configured else varies by platform when not configured.	Numeric, such as 297, not changeable.
&FOCLANGCODE	Language of the server when NLS is configured, else blank.	String, such as FRE, not changeable.
&LINES	Number of lines printed in last TABLE.	Numeric, not changeable.
&RECORDS	Number of records retrieved for last TABLE/GRAPH.	Numeric, not changeable.
&TRANS	Number of transactions processed.	Numeric, not changeable.

System Variable	Description	Format or Value
&ACCEPTS	Accepted transactions.	Numeric, not changeable.
&NOMATCH	Nomatch rejects.	Numeric, not changeable.
&FORMAT	Format errors.	Numeric, not changeable.
&INVALID	Invalid conditions.	Numeric, not changeable.
&DUPLS	Duplicates rejected.	Numeric, not changeable.
&INPUT	Segments input.	Numeric, not changeable.
&CHNGD	Values updated.	Numeric, not changeable.
&DELTED	Segments deleted.	Numeric, not changeable.
&RETCODE	Value after DOS command.	Numeric, not changeable.
&BASEIO	Number of input/output operations.	Numeric, not changeable.
&READS	Read operations from external file.	Numeric, not changeable.
&REJECTS	Number of rejected transactions.	Numeric, not changeable.

Example: Using Predefined Run-Time System Variables

The following example incorporates the system variable &DATE into an SQL request, testing a user-supplied variable (IDATE) against it.

```
SQL
SELECT '&DATE',IDATE
FROM filename
WHERE IDATE < '08/08/2004'
.
.
.
```


Supplying Values for Variables

You must supply values for variables in a procedure even if the value is a blank. For instance, some server commands are invalid without values but process normally with blanks.

Supply values for variables in the following ways:

- With the EXEC command as parameters.
- With a command such as -DEFAULTS, -SET, or -READ.

This section describes these methods.

General Rules

The following general rules apply to values for variables:

- The maximum length is 32,000 characters when using the -TYPE or -WRITE command.
- Once a value is supplied for a local variable, it is used throughout the procedure unless it is changed with a command such as -SET or -READ.
- Once a value is supplied for a global variable, it is used throughout the session in all procedures unless it is changed with -SET, -READ, or another command.

Supplying Values in the EXEC Command

The command EXEC enables you to call one procedure from another and set values for variables in the called procedure, using:

- Keyword parameters.
- Positional parameters.
- A combination of keyword and positional parameters.

Syntax: How to Pass Keyword Parameters

```
EX[EC] procedurename=value[, ...]
```

where:

procedure

Is the name of the called procedure.

name=value

Is a keyword parameter.

If *value* contains an embedded comma, blank, or equal sign, it must be enclosed in single quotation marks. For example:

```
EX SLRPT AREA=S, CITY='NY, NY'
```

Name=value pairs must be separated by commas. You do not need to enter pairs in the order in which they are encountered in the procedure.

A procedure with a long list of variables or long values for parameters may be broken onto multiple lines by inserting a comma as the last character on the line and entering the rest of the list on the following line.

Syntax: **How to Pass Positional Parameters**

```
EX[EC] procedure parm1[, ...]
```

where:

procedure

Is the name of the called procedure.

parm1

Is a positional parameter. You do *not* need to specify the number in the parameter list. Dialogue Manager matches the values, one by one, to the positional variables as they are encountered in the called procedure.

However, you must specify the parameters in the order in which the numeric ampers are used in the procedure.

Consider the following procedure:

```
SQL
SELECT SUM(UNIT_SOLD),SUM(RETURNS),RETURNS/UNIT_SOLD FROM SALES
WHERE PROD_CODE BETWEEN '&1' AND '&2' AND CITY = '&3' ;
END
```

The calling procedure would be:

```
EX SLRPT B10,B20,STAMFORD
```

Reference: **Combining Positional and Keyword Parameters**

Consider a procedure with:

```
SQL
SELECT &1, &2, &FIELD1, &3 FROM CAR;
END
```

If it is called from another procedure using EXEC, the EXEC command and parameters might look like:

```
EXEC PPARAM1 MODEL,MPG,FIELD1=CAR,COUNTRY
```

Displaying Execution Flow

Dialogue Manager implements IF THEN ELSE and other flow logic such as -GOTO. The flow logic and other syntax may be displayed (typically for debugging purposes) by using the &ECHO variable.

How to Display Command Lines While Executing

```
EX xxx ECHO=value
```

or

```
-DEFAULTS &ECHO = value
```

or

```
-SET &ECHO = value;
```

Valid values are:

ON

Displays lines that are expanded and stacked for execution.

ALL

Displays Dialogue Manager commands as well as lines that are expanded and stacked for execution.

OFF

Suppresses display of both stacked lines and Dialogue Manager commands. OFF is the default value.

Note that if the procedure is encrypted, &ECHO automatically receives the value OFF, regardless of the value that is assigned explicitly.

-DEFAULT[S] Command

The Dialogue Manager command -DEFAULTS supplies an initial (default) value for a variable that had no value before the command was processed. It ensures that values are provided for variables whether or not they are provided elsewhere.

How to Supply Values With the -DEFAULT[S] Command

```
-DEFAULT[S] &[&]name=value [,...]
```

where:

&name

Is the name of the variable.

value

Is the default value assigned to the variable. Enclose in single quotation marks (') if there are embedded spaces. Use a comma (,) as a separator when there are multiple variables on a line. When a single variable per line is used, use of an "ending" semicolon (;) is allowed, but not required. The command may be repeated as many times as needed (it is not necessary to use the multiple variable form to set multiple variables).

-SET Command

The -SET command is used to assign a value or computed value to a variable.

How to Supply Values With the -SET Command

`-SET &[&]name={expression|value};`

where:

&name

Is the name of the variable. A double ampersand indicates that the variable is global.

expression

Is a valid expression. An expression may occupy several lines. Each continuation line requires a leading dash (-) in the first character position to indicate that it is a continuation. A -SET always ends with a semi-colon (;).

value

Is a value such as a number or string (enclose the string in single quotation marks (') if contains an embedded space).

The following are some examples:

Example (single line)	Purpose
<code>-SET &NUM = 1 ;</code>	Set value.
<code>-SET &FN = JOHN ;</code>	Set value.

Example (single line)	Purpose
<code>-SET &LN = 'JOHN' ;</code>	Set value with single quotation marks (').
<code>-SET &NAME = 'JOHN DOE' ;</code>	Set value with single quotation marks (') required by embedded space.
<code>-SET &NAME = &FN ' ' &LN ;</code>	Set value as concatenated string.
<code>-SET &NAME = &LN ', ' &LN ;</code>	Set value using hard concatenated string.
<code>-SET &NAME = IF &NAME EQ 'JOHN DOE' THEN 'JOHNATHAN DOE' ELSE &NAME ;</code>	Set value using expression.

Branching

The execution flow of a procedure is determined using the following commands:

- `-GOTO`. Used for unconditional branching, `-GOTO` transfers control to a label.
- `-IF...GOTO`. Used for conditional branching, `-IF...GOTO` transfers control to a label depending on the outcome of a test.

Syntax: How to Use the `-GOTO` Command for Unconditional Branching

```
-GOTO label
.
.
.
-label [TYPE text]
```

where:

label

Is a user-defined name of up to 64 characters. Embedded blanks in labels or using other Dialogue Manager commands as labels is not allowed. It is also advised to not use words that may be confused with functions or arithmetic or logical operations. If a label is not found by searching downward in the procedure, it is searched from the beginning a single time (if still not found a user error will occur).

`TYPE text`

Optionally, sends a message to the application.

Syntax: **How to Use the -IF...GOTO Command for Conditional Branching**

```
-IF expression [THEN] GOTO label ;  
-IF expression [THEN] GOTO label ELSE GOTO label ;  
-IF expression [THEN] GOTO label ELSE IF... ;  
. ;  
. ;  
-label [TYPE text]
```

where:

label

Is a user-defined name of up to 64 characters. Embedded blanks in labels, or using other Dialogue Manager commands as labels is not allowed. It is also advised to not use words that may be confused with functions or arithmetic or logical operations. If a label is not found by searching downward in the procedure, it is searched from the beginning a single time (if still not found a user error will occur).

expression

Is a valid expression. Literals need not be enclosed in single quotation marks (') unless they contain embedded blanks or commas (,).

THEN

Is an optional keyword that increases readability of the command.

ELSE GOTO

Optionally, passes control to *label2* when the -IF test fails.

ELSE IF

Optionally, specifies a compound -IF test.

TYPE *text*

Optionally, sends a message to a client application.

Continuation lines must begin with a hyphen (-) and lines must break between words. A space after the hyphen is not required, but it adds to readability.

Screening Values With -IF Tests

To ensure that a supplied value is valid in a procedure, test for its:

- Presence
- Type
- Length

For instance, you would not want to perform a numerical computation on a variable for which alphanumeric data has been supplied.

Syntax: How to Test for the Presence of a Value

```
-IF &name.EXIST GOTO label...;
```

where:

&name

Is a user-supplied variable.

`.EXIST`

Indicates that you are testing for the presence of a value. If a value is not present, a zero (0) is passed to the expression. Otherwise, a non-zero value is passed.

`GOTO label`

Specifies a label to branch to.

Syntax: How to Test for the Length of a Value

```
-IF &name.LENGTH expression GOTO label...;
```

where:

&name

Is a user-supplied variable.

`.LENGTH`

Indicates that you are testing for the length of a value. If a value is not present, a zero (0) is passed to the expression. Otherwise, the number of characters in the value is passed.

expression

Is the remainder of a valid expression, such as `GT 8`.

GOTO label

Specifies a label to branch to.

Syntax: **How to Test for the Type of a Value**

-IF &name.TYPE expression GOTO label...;

where:

&name

Is a user-supplied variable.

.TYPE

Indicates that you are testing for the type of a value. The letter N (numeric) is passed to the expression if the value is interpreted as a number up to 10^9-1 and is stored in four bytes as a floating-point format. In Dialogue Manager, the result of an arithmetic operation with numeric fields is truncated to an integer after the whole result of an expression is calculated. If the value could not be interpreted as numeric, the letter A (alphanumeric) is passed to the expression.

expression

Is the remainder of a valid expression, such as EQ A.

GOTO label

Specifies a label to branch to.

Calling Another Procedure

One procedure calls another procedure using:

- The `-INCLUDE file` command, which incorporates a whole or partial procedure into the current procedure and executes it immediately when encountered. Unlike EXEC, no parameters may be passed with `-INCLUDE`, but they may be `-SET` on lines that are prior to the `-INCLUDE`. A `-INCLUDE` procedure can be a partial procedure that may contain header text or code to include at run time based on a step in the originating procedure.
- The command EXEC. The command is stacked and executed when the appropriate Dialogue Manager command is encountered. The called procedure must be fully executable.

Syntax: **How to Use the -INCLUDE Command**

Lines incorporated with a `-INCLUDE` are processed as though they had been placed in the calling procedure originally.


```
-INCLUDE filename
```

where:

filename

Is the name of the called procedure.

A calling procedure cannot branch to a label in a called procedure, and vice versa.

Nesting

Any number of different procedures may be invoked as a `-INCLUDE` from within a single calling procedure. Called procedures can also call (nest) additional procedures. The same procedure may even call itself recursively. However, recursive calls are limited to four levels deep.

The EXEC Command

A procedure also calls another one with the command `EXEC`. The called procedure must be fully executable.

See [Supplying Values for Variables](#) on page 509 for additional information about passing parameter values.

Use the EXEC (or EX) Command

Procedure lines with `EXEC` (or `EX`) commands temporarily stop execution of the current procedure and call the new procedure. The new procedure may be passed current variables that are available as parameter values to the new procedure, but must do so explicitly unless those variables are global variables.

```
EXEC [appname/]filename [parm[,...]]
```

where:

appname

Is the location of the called procedure.

filename

Is the name of the called procedure.

parm[,...]

Are one or more parameters being passed, either positionally or by name.

ON TABLE HOLD

When a server receives the results of an SQL request (an answer set) from another server, the answer set will either:

- ❑ Be returned to the client application using ON TABLE PCHOLD. That command is described in *ON TABLE PCHOLD* on page 519.
- ❑ Be held on the initiating server, without sending it back to the client application, using ON TABLE HOLD. A corresponding Master File for the file that holds the answer set is also created.

Syntax: How to Use the ON TABLE HOLD Command

```
SQL
SQL request;
ON TABLE HOLD [AS filename] FORMAT format
END
```

where:

filename

Is the name of the file that holds the answer set. If *filename* is omitted, the name of the held file on the server is HOLD, and subsequent creations of HOLD files overlay each other. The file name is a symbolic name known to the operating system for the server environment.

format

A valid format option for the server. The following is a list of commonly used values are:

```
ALPHA, BINARY, COMMA, COMT, DB2, DIF, DOC (WebFOCUS ONLY),
EXCEL, EXL2K (WebFOCUS ONLY), EXL2K PIVOT (WebFOCUS ONLY),
FOCUS, HTML, HTMTABLE, INGRES, INTERNAL, LOTUS, PDF, POSTSCRIPT,
SQL, SQLDBC, SQLINF, SQLMAC, SQLMSS, SQLODBC, SQLORA,
SQLSYB, SYLK, TABT, WK1, and WP.
```

The above list is not intended to be a complete list, and the use of some of these formats is limited to use with a configured adapter.

END

Is required on a separate line.

ON TABLE PCHOLD

In order for a Dialogue Manager procedure to return an answer set to a client application, a certain set of commands must be issued directly after the SQL request in the syntax of the procedure.

Reference: Use the ON TABLE PCHOLD Command

```
SQL
SQL request;
TABLE
ON TABLE PCHOLD [FORMAT ALPHA]
END
```

where:

```
SQL request;
```

Is a valid SQL request, ending with a semicolon (;).

```
FORMAT ALPHA
```

Optionally specifies that the hold file on the client is a text file. Use any valid format available on the client, but the underlying transfer is in alpha format. FORMAT ALPHA is the default value.

```
END
```

Is required on a separate line.

Example: Using the ON TABLE PCHOLD Command

This example shows how the ON TABLE PCHOLD command requests information from a table in a catalog.

```
SQL
SELECT NAME, CREATOR, COLCOUNT, RECLENGTH FROM SYSTABLE;
TABLE
ON TABLE PCHOLD FORMAT ALPHA
END
```

The result of the request is an answer set sent to the client application by the server.

Creating Expressions

An expression consists of variables and literals (numeric or alphanumeric constants) that are combined arithmetically, logically, or in some other way to create a new value.

This section describes how to create:

- ❑ [Arithmetic Expressions](#) on page 520.
- ❑ [Alphanumeric Expressions](#) on page 521.
- ❑ [Logical Expressions](#) on page 522.
- ❑ [Compound Expressions](#) on page 525.

Arithmetic Expressions

An arithmetic expression is:

- ❑ A numeric constant, for example, 1.
- ❑ Two variables joined by one of the following arithmetic operators:

Operator	Meaning
+	Addition
-	Subtraction
*	Multiplication
/	Division
**	Exponentiation

An example is:

```
&DELIVER_AMT / &OPENING_AMT
```

- ❑ Two or more arithmetic expressions, joined by one of the operators in the preceding list. An example is:

```
(&RATIO - 1) ** 2
```

- ❑ A compound expression or function that gives an arithmetic result.

Example: Using Arithmetic Expressions

Following are some arithmetic expressions used in the command -SET:

```
-SET &COUNT = 1;
-SET &NEWVAL = (&RATIO - 1) ** 2;
-SET &RATIO = (&DELIVER_AMT * 100) / (&OPENING_AMT);
```

Reference: Guidelines for Using Arithmetic Expressions

Keep the following in mind as you create arithmetic expressions:

- If you attempt to divide by 0, Dialogue Manager sets the result to 0.

Arithmetic operations are performed before logical operations, in the following order:

Operator	Meaning
**	Exponentiation
/ *	Division and multiplication
+ -	Addition and subtraction

- For operations on the same level (for example, division and multiplication), the evaluation is performed from left to right.
- An expression in parentheses is evaluated before any other expression.
- Values for local and global variables (amper variables) are stored internally as character strings, including numeric values. If a calculation is performed on an amper variable, the variable is first converted from a character string into a numeric. After the whole result is calculated, the result of arithmetic operations with numeric fields is truncated to the integer field. Finally, the result is converted back into a character string.

Alphanumeric Expressions

An alphanumeric expression is:

- A literal enclosed in single quotation marks, for example, 'Smith John'.
- A logical expression that yields an alphanumeric result.
- A function that yields an alphanumeric result.
- Two or more alphanumeric variables or literals combined into a single string. See [How to Concatenate Alphanumeric Variables and Literals](#) on page 522 for the syntax and an example.

Syntax: **How to Concatenate Alphanumeric Variables and Literals**

```
variablename = {alphaexp1|'literal'} concatenation
{alphaexp2|'literal'} [...]
```

where:

variablename

Is the name of the variable assigned to the result of the concatenation.

alphaexp1, alphaexp2

Are local or global variable that forms part of the concatenation.

literal

Is a literal that forms part of the concatenation. It must be enclosed in single quotation marks (').

concatenation

Is one of the following symbols:

	Indicates strong (also known as hard) concatenation, which suppresses trailing blanks.
	Indicates weak concatenation, which preserves individual field lengths, including trailing blanks.

Example: **Concatenating Alphanumeric Variables and Literals**

```
-SET &NAME = &LASTNAME || ', ' || &FIRST_INIT;
```

If &LASTNAME is equal to Doe and &FIRST_INIT is equal to J, &NAME is set to:

Doe,J

Logical Expressions

A logical expression contains logical and relational operators and is evaluated to a value that is true or false.

Example: Forming a Logical Expression

This example shows various elements that are used to form a logical expression. The abbreviation *exp* stands for expression.

```
{arithmetic exp|alphanumeric exp} operator1
{numeric lit|alphanumeric lit} OR...
expression operator2 expression
logical exp {AND|OR} logical exp
NOT logical exp
```

where:

operator1

Is one of the following: EQ, NE, OMITS, or CONTAINS.

expression

Is either an arithmetic, alphanumeric, or logical expression.

operator2

Is one of the following: EQ, NE, LE, LT, GE, or GT.

The following table defines valid operators (EQ, NE, and so on) used in this example.

Operator	Description
EQ	Tests for a value equal to another value.
NE	Tests for a value not equal to another value.
OMITS	Tests for a value that does not contain a matching character string.
CONTAINS	Tests for a value that does contain a matching character string.
LE	Tests for a value less than or equal to another value.
LT	Tests for a value less than another value.
GE	Tests for a value greater than or equal to another value.
GT	Tests for a value greater than another value.

Operator	Description
AND	Returns a value of true if both of its operands are true.
OR	Returns a value of true if either of its operands is true.
NOT	Returns a value of true if the operand is false.

Reference: Guidelines for Alphanumeric and Logical Expressions

Keep the following in mind:

- ❑ An alphanumeric literal with embedded blanks or commas must be enclosed in single quotation marks. For example:

```
-IF &NAME EQ 'JOHN DOE' GOTO QUIT;
```

To produce a single quotation mark within a literal, place two single quotation marks where you want one to appear:

```
-IF &NAME EQ 'JOHN O' 'HARA' GOTO QUIT;
```

- ❑ A computational field may be assigned a value by equating it to a logical expression. If the expression is true, the field has a value of 1; if the expression is false, the field has a value of 0.
- ❑ Use OR to connect literals or other expressions. You must also use parentheses to separate expressions connected with OR.
- ❑ Logical operations are done after arithmetic operations, in the following order:

```
EQ NE LE LT GE GT NOT CONTAINS OMTS
```

```
AND
```

```
OR
```

- ❑ Separate a collection of test values with OR:

```
-IF &STATE EQ 'NY' OR 'NJ' OR 'WA' GOTO QUIT;
```

In this case, OR and EQ are evaluated at the same level.

- ❑ Use parentheses to specify a desired order. An expression in parentheses is evaluated before any other expression. For example, the command


```
-IF &STATE EQ 'NY' AND &COUNTRY EQ 'US' OR 'UK' THEN...
```

is evaluated as:

```
IF &STATE EQ 'NY' IF &COUNTRY EQ 'US'...
```

Dialogue Manager then evaluates the phrase OR UK and indicates that it is a syntax error.

To write the command correctly, add parentheses:

```
-IF ((&STATE EQ 'NY') AND (&COUNTRY EQ 'US' OR 'UK')) THEN...
```

Compound Expressions

A compound expression has the following form:

```
-IF expression THEN expression ELSE expression;
```

The following restrictions apply:

- Each of the expressions may itself be a compound expression, although the expression following -IF may not be a -IF...THEN...ELSE expression (for example, -IF...-IF...).
- If the expression following THEN is itself a compound expression, it must be enclosed in parentheses; this rule does not apply to an expression following ELSE.
- Compound expressions only have up to 16 -IF commands.

Example: Using Compound Expressions

If the following example is executed without an input parameter list, the client application receives the message NONE. If it executes with the parameter BANK='FIRST NATIONAL', the client application receives the message FIRST NATIONAL.

```
-DEFAULTS &BANK = ' '
-SET &BANK = IF &BANK EQ ' ' THEN 'NONE'
-ELSE &BANK;
-TYPE &BANK
```

The next example uses a compound expression to define a truth condition (1 is true and 0 is false).

```
-DEFAULTS &CURR_SAL = 900,&DEPARTMENT=MIS
-SET &MYTEST = (&CURR_SAL GE 1000) OR (&DEPARTMENT EQ MIS);
-IF &MYTEST EQ 1 THEN GOTO YES ELSE GOTO NO;
-YES
-   TYPE YES
-EXIT
-NO
-   TYPE NO
```

When this code is executed, &MYTEST evaluates before the -IF, and the client application receives the message YES, unless alternate values causing an untrue condition are passed on execution.

ON TABLE HOLD

When a server receives the results of an SQL request (an answer set) from another server, the answer set will either:

- Be returned to the client application using ON TABLE PCHOLD. That command is described in *ON TABLE PCHOLD* on page 519.
- Be held on the initiating server, without sending it back to the client application, using ON TABLE HOLD. A corresponding Master File for the file that holds the answer set is also created.

Syntax: How to Use the ON TABLE HOLD Command

```
SQL
SQL request;
ON TABLE HOLD [AS filename] FORMAT format
END
```

where:

filename

Is the name of the file that holds the answer set. If *filename* is omitted, the name of the held file on the server is HOLD, and subsequent creations of HOLD files overlay each other. The file name is a symbolic name known to the operating system for the server environment.

format

A valid format option for the server. The following is a list of commonly used values are:

```
ALPHA, BINARY, COMMA, COMT, DB2, DIF, DOC (WebFOCUS ONLY),
EXCEL, EXL2K (WebFOCUS ONLY), EXL2K PIVOT (WebFOCUS ONLY),
FOCUS, HTML, HTMTABLE, INGRES, INTERNAL, LOTUS, PDF, POSTSCRIPT,
SQL, SQLDBC, SQLINF, SQLMAC, SQLMSS, SQLODBC, SQLORA,
SQLSYB, SYLK, TABT, WK1, and WP.
```

The above list is not intended to be a complete list, and the use of some of these formats is limited to use with a configured adapter.

END

Is required on a separate line.

ON TABLE PCHOLD

In order for a Dialogue Manager procedure to return an answer set to a client application, a certain set of commands must be issued directly after the SQL request in the syntax of the procedure.

Reference: Use the ON TABLE PCHOLD Command

```
SQL
SQL request;
TABLE
ON TABLE PCHOLD [FORMAT ALPHA]
END
```

where:

SQL request;

Is a valid SQL request, ending with a semicolon (;).

FORMAT ALPHA

Optionally specifies that the hold file on the client is a text file. Use any valid format available on the client, but the underlying transfer is in alpha format. FORMAT ALPHA is the default value.

END

Is required on a separate line.

Example: Using the ON TABLE PCHOLD Command

This example shows how the ON TABLE PCHOLD command requests information from a table in a catalog.

```
SQL
SELECT NAME, CREATOR, COLCOUNT, RECLENGTH FROM SYSTABLE;
TABLE
ON TABLE PCHOLD FORMAT ALPHA
END
```

The result of the request is an answer set sent to the client application by the server.

Platform-Specific Commands and Features

This topic describes platform-specific commands that may be included in a procedure for various purposes. However, it is always best to avoid the use of platform-specific commands, as they can affect the portability of an application to new environments.

The nature of z/OS, PDS Deployment, and security requires operating system commands to be a limited command set, and use of the keyword DYNAM to preface commands, where other environments are free to simply jump to the operating system command level and executable anything the application user has access to. DYNAM commands are also allowed in z/OS HFS Deployment.

On platforms other than z/OS PDS Deployment, native commands are directly available using the -SYSTEM, -UNIX, -WINNT, and -AS/400 series of commands.

File references that are symbolic logical names for the purpose of -READ, -WRITE, TABLE (using external files), and HOLD AS require a DYNAM ALLOC (PDS) or FILEDEF to create the logical reference.

For more information about reporting, see the *Creating Reports With TIBCO WebFOCUS® Language* manual. For more information about Dialogue Manager, see the *TIBCO WebFOCUS® Developing Reporting Applications* manual.

DYNAM Command (z/OS PDS and HFS Deployments)

The DYNAM command manipulates data sets under MVS.

```
DYNAM subcommand operand [operand]...
```

where:

subcommand

Is required, and specifies one of the operations (subcommands) in the list below. The abbreviated form of the subcommand's syntax is given under the full form. Details on each subcommand are provided in the following sections as noted.

ALLOCATE ALLOC ALLO	Allocates a data set. See <i>The ALLOCATE Subcommand</i> on page 531.
CONCAT CONC	Concatenates data sets. See <i>The CONCAT Subcommand</i> on page 540.
FREE	Frees data sets specified by ddnames or dsnames. Names may contain wildcard characters. See <i>The FREE Subcommand</i> on page 541.
CLOSE CLO	Closes data sets. Use this subcommand when data sets cannot be freed because of being open. See <i>The CLOSE Subcommand</i> on page 542.
COPY	Copies an entire data set or selected partitioned data set (PDS) members. This subcommand provides features such as record format conversion, either automatic or option controlled. See <i>The COPY Subcommand</i> on page 543.
COPYDD	Copies a sequential data set or PDS member. COPY handles all the features of COPYDD, and is recommended for use instead of COPYDD. See <i>The COPYDD Subcommand</i> on page 545.
DELETE DEL	Deletes an entire data set or selected PDS members. See <i>The DELETE Subcommand</i> on page 546.
RENAME REN	Renames an entire data set or selected PDS members. See <i>The RENAME Subcommand</i> on page 547.
SUBMIT SUB	Submits MVS jobs. See <i>The SUBMIT Subcommand</i> on page 548.
COMPRESS COMP	Compresses a PDS. See <i>The COMPRESS Subcommand</i> on page 549.

operand

May be a keyword, a keyword followed by its parameter, or a parameter without a keyword.

The following rules apply to the DYNAM command:

- ❑ The subcommand, keywords, and parameters are separated with one or more blanks. Keywords are coded in free format.

- ❑ A parameter may be a list of subparameters (for example, VOLUME for a multi-volume data set). Separate subparameters in the list using commas. To include blanks between subparameters (with or without the comma), enclose the entire list in parentheses. For example:

```
A,B (A,B) (A B) (A, B) (A,B C, D)
```
- ❑ A DYNAM command may span several lines. Enter a hyphen (-) at the end of each line to be continued. When the lines are concatenated, blanks after the hyphen and leading blanks from the next line are removed. Blanks before the hyphen are removed if they are preceded by a comma. The total length of a DYNAM command may not exceed 2,048 characters.
- ❑ Most keywords may be truncated up to the shortest unambiguous length. The commonly used abbreviations are fixed. Note that the unique truncation of a keyword may not always be valid as new keywords are added. It is recommended that the full keyword be used in files.
- ❑ Fixed abbreviations are given in the following sections on the subcommands. For example, DDNAME may be abbreviated as DD, DDN, DDNA, DDNAM, or DDNAME.
- ❑ Certain keywords have synonyms. For example, the keywords FILENAME and DDNAME are synonyms, and so are DATASET and DSNAME.
- ❑ As in TSO, a data set name is enclosed in single quotation marks. Prefix substitution is not supported; specify only the fully qualified data set names in JCL.
- ❑ Some DYNAM commands accept either the ddname or data set name (dsname) as the same parameter. In such cases, the parameter is considered a ddname if it is not longer than 8 bytes, does not contain periods (.), and is not enclosed in single quotation marks (''). Otherwise, the parameter is considered a data set name. Thus, to specify an unqualified data set name, enclose it in single quotation marks.

Use of Data Sets

z/OS PDS Deployment obtains a lock for any allocated data set name; a shared lock for those specified as SHR, and an exclusive lock for OLD, NEW, or MOD.

Although data sets are allocated more than once in a job step, only one type of lock may be obtained. For example, if the data set is initially allocated as SHR and is then allocated as OLD in the same step, the MVS lock changes from shared to exclusive, and the data set is not available for use by other jobs until *all* allocations in the job are freed.

The DYNAM commands that manipulate data sets use an improved locking mechanism, similar to that implemented in ISPF:

- ❑ Any output PDS is allocated by DYNAM (or pre-allocated by the user) as SHR. This avoids exclusive MVS locking, which lasts until all data set allocations are freed.
- ❑ To protect from simultaneous updating, DYNAM obtains an exclusive lock as used by ISPF (and other programs, including LINKEDIT), but only during the actual update time. This lock controls access to the data set between users, even from within the same job step.

Note: The DYNAM locking mechanism protects from simultaneous updating and possible corruption of data, but does not protect from updating and simultaneous reading. For example, it is possible to continue to read a PDS member recently deleted by another user.

The ALLOCATE Subcommand

The DYNAM ALLOCATE command allocates a data set.

Using the ALLOCATE Subcommand

```
DYNAM ALLOCATE [disposition] [CLOSE]

DDNAME ddname [DEFER] [DSNAME dsname[(memname)]] [DUMMY]
[EXPDT date]
[HIPER OFF]
[DALDKYL labelname]

[INPT|OUTPT]

[LABEL type]
[MEMBER memname] [status] [MSVGP msvgp]

[PARALLEL] [PASSWORD password] [PERM] [POSITION nnnn]
[REFVOL dsname] [RETPD days] [REUSE]
[UNIT unit]
[VOLUME volser]
```

Space operands are:

```
[format]

[parameter]

[DIR n]
[PRIMARY n1]
[RELEASE] [ROUND]
[SECONDARY n2] [SPACE space]
```

DCB operands are:

```
[BLKSIZE n] [BUFNO n]  
[DEN n] [DSORG dsorg]  
[LRECL n]  
[EATTR {NO|OPT}]  
[RECFM recfm] [REFDD ddname] [REFDSN dsname]
```

SMS and VSAM operands are:

```
[DATACLASS name] [DSNTYPE {LIBRARY|PDS|EXTREQ|EXTPREF}]  
  
[KEYOFF n]  
[LIKE dsname]  
[MGMTCLASS name]  
[RECORD recorg]  
[SECMODEL name] [STORCLASS name]  
[BUFND m]  
[BUFNI n]
```

Output printing operands are:

```
[DEST dest[.user]]  
  
[FCB name [ALIGN|VERIFY]] [FORMS name]  
  
[HOLD]  
[OUTLIM n] [OUTPUT name]  
[SYSOUT class]  
[USER user]  
[WRITER name]
```

where:

ALLOCATE

Can be abbreviated as ALLOC or ALLO.

disposition

Is one of the following:

```
CATALOG  
DELETE  
KEEP  
UNCATALOG
```

By default, for a data set status of NEW, if *dsname* is specified, the disposition is CATALOG; otherwise, the disposition is DELETE. Is incompatible with SYSOUT. CATLG and UNCAT are also valid as synonyms for CATALOG and UNCATALOG. DELETE, KEEP, and UNCATLOG follow the standard MVS meanings of delete after free, keep as is after free, and keep uncataloged.

CLOSE

Is deallocation of the data set at close, rather than at the end of the step. The JCL analogy is FREE=CLOSE.

DDNAME *ddname* DD

Is the DDNAME to be associated with an allocation; it must be specified. Synonym is FILENAME.

DEFER

Assigns device(s) to the data set but defers mounting of the volume(s) until the data set is opened. The JCL analogy is DEFER in UNIT.

DSNAME *dsname* [(*memname*)]

The member name is specified either in parentheses after *dsname* or using keyword MEMBER (see also MEMBER). If *dsname* is specified as an asterisk (*), terminal is allocated. This is used for output only. Synonym is DATASET.

DUMMY

Allocates a dummy data set.

DALDKYL *labelname*

Specifies the label for the encryption key used by the system to encrypt the data set. The maximum length of the label is 64 characters. The TSO ALLOC equivalent is DSKEYLBL.

For example, to allocate a DD with pervasive encryption, key MYLABELNAME:

```
DYNAM ALLOC DD ASDF . . . . DALDKYL MYLABELNAME . . . SHR
```

This keyword has an effect only when creating an extended format data set. the first value for the DISP keyword must be coded or defaulted.

EXPDT *date*

Is the expiration date in format YYDDD, YYYY/DDD, or YYYYDDD. Is incompatible with RETPD and SYSOUT.

HIPER OFF

Prohibits allocation in a hiperspace. Is equivalent to UNIT NOHIPER, and is used when UNIT is also to be specified. For example, UNIT VIO HIPER OFF.

INPT

Data set is to be processed as input only (INPT) or output only (OUTPUT). The JCL analogy is IN in LABEL. Is incompatible with SYSOUT.

OUTPT

Data set is to be processed as input only (INPT) or output only (OUTPUT). JCL analogy: IN in LABEL. Is incompatible with SYSOUT.

LABEL *type*

Specifies type of volume labels. Can be one of the following: NL, SL, NSL, SUL, BLP, LTM, AL, or AUL. Is incompatible with SYSOUT.

MEMBER *memname*

Is the name of a PDS member to be allocated. See also DSNAME.

status

Is the data set status. Possible values are:

NEW is the default data set status. Incompatible with SYSOUT.

MOD is an extended data set.

OLD is exclusive control of the data set.

SHR is shared access to the data set.

MSVGP *msvgp*

Is the identification of a group of mass storage system (MSS) virtual volumes. Is incompatible with SYSOUT and VOLUME.

PARALLEL

Each volume is to be mounted on a separate device. The JCL analogy is P in UNIT.

PASSWORD *password*

Password for a password-protected data set.

PERM

The allocation is to be permanent—that is, protected from being freed or concatenated by any DYNAM command issued by an MSO user. The operand is valid only in an MSO server initialization profile.

POSITION *nnnn*

Data set sequence number on a tape volume, up to 9999. The JCL analogy is the first subparameter in LABEL.

REFVOL *dsname*

Volume serial information is to be obtained from the named cataloged data set. The JCL analogy is VOL=REF=*dsname*. Is incompatible with SYSOUT and VOLUME.

RETPD days

Is the retention period, up to 9999 days. Is incompatible with EXPDT and SYSOUT.

REUSE REU

If the ddname to be allocated is already in use, it is to be freed.

UNIT unit

Is the device group name, device type, specific unit address, or NOHIPER. NOHIPER prohibits allocation in a hiperspace, and is meaningful for a temporary (NEW, DELETE) data set; see also HIPER OFF.

VOLUME volserVOL

Are volume serial numbers. Are incompatible with REFVOL and SYSOUT. Synonyms are VOLume and VOLser.

Space operands may be:

format

The format of the primary space to be allocated. Possible values are:

ALX is up to five contiguous areas.

CONTIG is one contiguous area.

MXIG is one maximal contiguous area.

JCL analogy: ALX/CONTIG/MXIG in SPACE.

n

Represents units of primary and secondary space allocation.

parameter

The parameter for space allocation. Possible values are:

BLOCKS [n]
CYLINDERS
MEGABYTES
PAGES
TRACKS

n represents units of primary and secondary space allocation. If the parameter for BLOCKS is omitted, the average block length is copied from BLKSIZE. If the space unit is omitted but SPACE and BLKSIZE are specified, BLOCKS equal BLKSIZE is used. For PAGES, BLOCKS 4096 is used. BLKSIZE must be specified if the BLOCKS parameter is specified.

Synonyms are CYLs for CYLINDERS and TRKs for TRACKS.

DIR *n*

The number of 256-byte records for the directory of a PDS.

PRIMARY *n1*

The primary space quantity. See also SPACE.

RELEASE

The unused space is to be released when the data set is closed. Synonym is RLSE.

ROUND

If space is requested in BLOCKS, MEGABYTES, or PAGES, it is to be rounded to whole cylinder(s).

SECONDARY *n2*

Is the secondary space quantity. See also SPACE.

SPACE *space SP*

The primary (n1) and/or secondary (n2) space quantity in one of the following formats:

n1/(n1)/n1,n2/(n1,n2)/n1 n2/(n1 n2)/,n2/(,n2)

See also PRIMARY and SECONDARY.

DCB operands may be:

BLKSIZE *n*

The block size, up to 32760. See also BLOCKS.

BUFNO *n*

The number of buffers, up to 255.

DEN *n*

n represents magnetic tape density: 0, 1, 2, 3, or 4 for 200, 556, 800, 1600, 6250 bpi respectively.

DSORG *dsorg*

The data set organization. Default, for NEW only: PO if DIR or DSNTYPE specified; PS otherwise. The following values are syntactically correct:

VS	VSAM
----	------

<code>PO/POU</code>	PDS or PDS unmovable.
<code>DA/DAU</code>	Direct access or direct access unmovable.
<code>PS/PSU</code>	Physical sequential or physical sequential unmovable.

`LRECL n`

The logical record length, up to 32760.

`EATTR {NO|OPT}`

Specifies if the data set to be created on a DASD volume can have extended attributes and reside in the extended address space (cylinder-managed space on the volume beyond the 64k cylinder limit). Valid values are:

- NO.** Data set to be created is not allowed to have extended attributes. It must not be created in the extended address space. This is the default value.
- OPT.** Data set to be created allows extended attributes to be stored in a 'FORMAT 8/9 DSCB' on the volume. The data set is allowed to be created in the extended address space. For example:

```
DYNAM ALLOC DD EAV DSN hlq.EAV.DATA NEW DSORG PS RECFM FB LRECL 80
CYL EATTR OPT
```

`RECFM recfm`

The record format. The first letter must be D, F, U, or V, which may be followed by any valid combination of A, B, M, S, or T:

<code>A</code>	Records with ISO/ANSI control characters.
<code>B</code>	Blocked records.
<code>D</code>	Variable-length ISO/ANSI tape records.
<code>F</code>	Fixed-length records.
<code>M</code>	Records with machine code control characters.
<code>S</code>	Standard fixed-length or spanned variable-length records.

T	Track overflow.
U	Undefined-length records.
V	Variable-length records.

REFDD *ddname*

DCB attributes are to be copied from the specified *ddname*. Under TSO, EXPDT and INPT/OUTPT specifications are also copied. Any of those can be overridden by the appropriate keyword on the same command. The JCL analogy is DCB=*.*ddname*. Is incompatible with REFDSN.

REFDSN *dsname*

DCB attributes (DSORG, RECFM, OPTCD, BLKSIZE, LRECL, RKP, KEYLEN) and EXPDT are to be copied from the specified cataloged data set. Any of those can be overridden by the appropriate keyword on the same command. The JCL analogy is DCB=*dsname*. Is incompatible with REFDD.

SMS and VSAM operands are:

DATACLASS *name*

The name of a data class for an SMS-managed data set.

DSNTYPE {**LIBRARY** | **PDS** | **EXTREQ** | **EXTPREF**}

LIBRARY is for a new partitioned extended (PDSE), and **PDS** is for a new partitioned data set. A PDSE cannot contain load modules, should be SMS-managed, and allows concurrent updating of different members.

EXTREQ specifies that a data set must be extended format (required). It may be VSAM or sequential.

EXTPREF requests that the data set be extended format (preferred). If the system cannot create it in extended format, it will try basic format. It may be VSAM or sequential.

KEYOFF *n*

The offset of the key in each logical record for a new VSAM key-sequenced (RECORD KS) data set.

LIKE *dsname*

Allocation attributes (DSORG, RECORD, or RECFM, LRECL, KEYLEN, KEYOFF, SPACE, DIR) are to be copied from the specified cataloged data set (model). Any of those can be overridden by the appropriate keyword on the same command.

MGMTCLASS name

The name of a management class for an SMS-managed data set.

RECOrg recorg

The VSAM record organization: KS, ES, RR, or LS for key-sequenced, entry-sequenced, relative record, or linear space data sets, respectively.

SECMODEL name

The data set RACF profile is to be copied from the named existing RACF profile.

STORCLASS name

The name of a storage class for an SMS-managed data set.

BUFND m

The number of VSAM DATA buffers.

BUFNI n

The number of VSAM INDEX buffers.

Output printing operands may be:

DEST dest[.user]

The remote destination for a SYSOUT data set. In conjunction with user ID, it is a node and a user at that node; the user ID is coded after the period (.) or using the USER keyword.

FCB name [ALIGN/VERIFY]

The name of an FCB (forms control buffer) image to be used for printing of a data set. The operator may be asked to check the printer forms alignment (ALIGN), or to verify the FCB image name displayed on the printer (VERIFY).

FORMS name FORM

A SYSOUT form name. JCL analogy: third subparameter in SYSOUT, FORMS in OUTPUT JCL.

HOLD

A SYSOUT data set is to be placed on the hold queue.

OUTLIM n

A limit for the number of logical records in a SYSOUT data set.

OUTPUT name

The name(s) of OUTPUT JCL statement(s) to be associated with a SYSOUT data set.

SYSOUT class

A SYSOUT data set is to be allocated and the specified output class (A-Z, 0-9) is to be assigned. If an asterisk (*) or NULL is coded, the class is copied either from CLASS in OUTPUT JCL if it is specified, or otherwise from MSGCLASS in JOB.

USER user

A SYSOUT data set is to be routed to the specified user ID. DEST is required to specify a user's node.

WRITER name

The name of an installation-written system output printing routine. The JCL analogy is the second subparameter in SYSOUT. Is incompatible with USER.

In addition to the shown fixed abbreviations and synonyms, keywords may be abbreviated up to the unique truncation. Those abbreviations are not fixed and may be changed when new keywords are added. They may be used interactively to save some keystrokes, but when a command is saved in a file, it is recommended that you use unabbreviated keywords.

Examples:

Allocate an existing data set:

```
DYNAM ALLOC DD MYDD DS MYID.DATA.SET SHR REU
```

Allocate a new data set. Defaults are NEW, CATALOG (*dsname* present), and DSORG PO (not-zero DIR present):

```
DYNAM ALLOC DD MYDD DS MYID.DATA.SET SPACE 6,2 TRACKS DIR 4 UNIT SYSDA -  
RECFM FB LRECL 80 BLKSIZE 1600
```

Allocate a terminal:

```
DYNAM ALLOC DD MYDD DS *
```

Allocate a SYSOUT data set with default output class. Upon freeing, the data set is sent to the user ID U1234 at node SYSVM:

```
DYNAM ALLOC DD MYDD SYSOUT * DEST SYSVM.U1234
```

The CONCAT Subcommand

The DYNAM CONCAT command concatenates up to 16 data sets.

```
DYNAM CONCAT [PERM] DDNAME ddname1 ddname2 [ddname3...]
```


where:

CONCAT

Can be abbreviated as CONC.

PERM

Is optional. This marks the concatenation as permanent—that is, protected from being freed or concatenated again by any DYNAM command issued by an MSO user. Valid only in an MSO server initialization profile.

DDNAME DDN DD

Are required; synonym is FILENAME.

ddname1

Is the first ddname to be concatenated and associated with the resulting concatenated group.

ddname2

Is the second ddname and any subsequent ddname to be concatenated.

For example:

```
DYNAM CONCAT DDN EDARPC MYEX NEWEX
```

The FREE Subcommand

The DYNAM FREE command deallocates any number of specified data sets.

```
DYNAM FREE {DDNAME ddname [ddname...] | DSNAME dsname [dsname...] }
```

where:

DDNAME DDNDD

Are required if there is no dsname; synonym is FILENAME.

ddname

Is the *ddname* of the data set to be freed.

DSNAME DSN DS

Are required if there is no ddname; the synonym is DATASET.

dsname

Is the name of the data set to be freed. All ddnames associated with this dsname, except concatenated groups, are deallocated.

While at least one ddname or data set name is required, you may specify more than one ddname or data set name. Each specified name may contain asterisks (*) and question marks (?) as wildcards. Wildcards are special characters used to specify a subset of names rather than one name. The wildcards appear anywhere in a name and mean the following:

* Represents any number of characters. For example, *Q* matches any name containing the character Q.

? Represents any single character. For example, ?Q? matches any 3-character name containing the character Q in the middle.

If the ddname is not found, a message is issued only if a single ddname without wildcards is specified. A message is not displayed if a data set or more than one ddname is not found.

Examples:

```
DYNAM FREE DDN SYS0* TEMP?
```

```
DYNAM FREE DSN MYID.DATA.SET
```

The CLOSE Subcommand

The DYNAM CLOSE command closes data sets that cannot be freed because they are opened.

```
DYNAM CLOSE {DDNAME ddname [ddname...] | DSNAME dsname [dsname...] }
```

where:

CLOSE

Can be abbreviated as CLO.

DDNAME DDN DD

Are required if there is no dsname; the synonym is FILENAME.

ddname

Is the ddname of the data set to be closed.

DSNAME DSN DS

Are required if there is no ddname; the synonym is DATASET.

dsname

Is the name of the data set to be closed. All ddnames associated with this dsname, except concatenated groups, are closed.

While at least one ddname or data set name is required, more than one ddname or data set name may be specified. Each specified name may contain wildcard characters. The same rules apply to the DYNAM CLOSE command as to the DYNAM FREE command.

The COPY Subcommand

The DYNAM COPY command copies an entire MVS data set or selected PDS members.

```
DYNAM COPY dname1 {[TO] dname2 [[MEMBER] members]|[MEMBER] members]}
[options]
```

where:

dname1

Is the dsname or ddname of the input data set. This is a positional parameter. It must precede all other operands.

TO

May be omitted if *dname2* does not match a reserved word, the MEMBER keyword, an option, or the TO keyword. To avoid confusion, use the TO keyword whenever *dname2* is a ddname.

dname2

Is the dsname or ddname of the output data set. If the output data set is not a PDS and the dsname is specified, it is allocated as OLD. If the ddname is specified, and the status is SHR, ensure that other users do not access the data set during COPY. Unlike ISPF, DYNAM locks a non-PDS data set in order to prevent simultaneous updating by different DYNAM users.

MEMBER

May be omitted if members are specified in parentheses.

members

Can be a single member specification or a list of member specifications. If the members are enclosed in parentheses, blanks before the left parenthesis may be omitted.

options

May be one or more of the following options:

APPEND adds the input to the end of the existing data, if the output is a sequential data set.

FORCE copies input DCB attributes (RECFM, BLKSIZE, LRECL, and KEYLEN) to the output data set. By default, only missing values are assigned.

KEYMOD allows key modification according to input/output **KEYLEN**: truncation or padding with binary zeros.

REPLACE replaces all output members matching the selected member names.

TRUNCATE allows truncation of input records that are longer than the output record length. Since trailing blanks are truncated automatically when **RECFM** is different, the keyword is used either to cut records of the same format or to cut non-blank data.

A member specification has the following syntax:

```
mem[ , [newmem] [ ,REPLACE ] ]
```

where:

mem

Is the selected member name.

newmem

Is the optional new name for the output member.

REPLACE

Is optional and specifies an existing member to be replaced in the output PDS.

Since the comma may be used in member specifications, they are separated with one or more blanks when specified in a list. Therefore, a list of member specifications is always enclosed in parentheses. For example:

```
(MEM MEM,NEWMEM MEM,NEWMEM,R MEM, ,R)
```

Note:

- All conversions between different DCB attributes (**RECFM**, **BLKSIZE**, and **LRECL**) are performed automatically.
- If the entire PDS is copied or any selected member's directory entry contains a **TTRN** in user data (for example, a load module), the IBM utility **IEBCOPY** is invoked. In this case, all options except **REPLACE** are ignored, format conversion is not possible, and copying members to the same PDS is not supported.
- If the main member and its alias names are copied, the relationship remains the same on the output PDS.
- If a specified **ddname** has been allocated with a member name, the data set is treated as sequential.

Examples:

Copies the entire data set, whether it is a PDS or not.

```
DYNAM COPY MYDD MYID.DATA.SET
```

All four commands are equivalent. Either input or output may be a sequential data set, or both are PDSs.

```
DYNAM COPY MYDD MYID.DATA.SET MEMBER MEM
DYNAM COPY MYDD MYID.DATA.SET (MEM)
DYNAM COPY MYDD(MEM) MYID.DATA.SET
DYNAM COPY MYDD MEMBER MEM MYID.DATA.SET
```

Copies and renames one member.

```
DYNAM COPY MYID.DATA.LIB TO MYDD(MEM1, MEM2)
```

Copies two members.

```
DYNAM COPY MYID.DATA.LIB TO MYDD(MEM1 MEM2)
```

Copies two members into same PDS with renaming.

```
DYNAM COPY MYDD(OLD1,NEW1,R OLD2,NEW2)
DYNAM COPY MYDD(OLD1,NEW1 OLD2,NEW2) REPL
```

The COPYDD Subcommand

The DYNAM COPYDD command copies a sequential data set or PDS member.

```
DYNAM COPYDD ddname1[ (mem1) ] ddname2[ (mem2) ]
```

where:

ddname1

Is the ddname of the input data set.

mem1

Is optional. It is the input member name.

ddname2

Is the ddname of the output data set.

mem2

Is optional. It is the output member name.

Note:

- ❑ If the specified ddname has been allocated with a member name, the data set is treated as sequential.
- ❑ Identically named members are always replaced on the output PDS.
- ❑ All conversions between different DCB attributes (RECFM, BLKSIZE, and LRECL) are performed automatically.
- ❑ Since the DYNAM COPY command has more features than COPYDD, it is recommended that you use COPY instead of COPYDD.

The DELETE Subcommand

The DYNAM DELETE command deletes an entire MVS data set or selected PDS members.

`DYNAM DELETE dsname`

where:

`DELETE`

Can be abbreviated as DEL.

`dsname`

Is the data set name to be deleted and uncataloged.

`dname`

Is the dsname or ddname of a PDS containing one or more members to be deleted. The ISPF-like lock is obtained.

`MEMBER`

May be omitted if the members are specified in parentheses.

`members`

Can be a single member name or a list of members. If the members are enclosed in parentheses, blanks before the left parenthesis can be omitted.

Examples:

```
DYNAM DELETE MYID.DATA.OLD
DYNAM DEL MYID.DATA.LIB MEMBER OLD1,OLD2
DYNAM DELETE MYDD(OLD1,OLD2)
DYNAM DEL MYDD(OLD1 OLD2 OLD3)
```

The RENAME Subcommand

The DYNAM RENAME command renames an entire MVS data set or selected PDS members.

```
DYNAM RENAME dsname1 dsname2
```

To rename individual members, use

```
DYNAM RENAME dname [MEMBER] members [REPLACE]
```

where:

RENAME

Can be abbreviated as REN.

dsname1

Is the data set name to be renamed and uncataloged.

dsname2

Is the new name to be assigned to the data set and cataloged.

dname

Is the dsname or ddname of a PDS containing one or more members to be renamed. The ISPF-like lock is obtained.

MEMBER

May be omitted if the members are specified in parentheses.

members

Can be a single member specification or a list of members. If the members are enclosed in parentheses, blanks before the left parenthesis can be omitted.

REPLACE

Is optional. This replaces all members matching the specified new names.

A member specification has the following syntax

```
oldmem,newmem[ ,REPLACE ]
```

where:

oldmem

Is the original member name.

newmem

Is the new member name.

REPLACE

Is optional and replaces existing members with the same name as *newmem*.

Since the comma is used in member specifications, each pair of members is separated with one or more blanks when specified in a list; therefore, a list of member specifications is always enclosed in parentheses.

Examples:

```
DYNAM RENAME MYID.DATA.OLD MYID.DATA.NEW
DYNAM REN MYID.DATA.LIB MEMBER OLD,NEW,R
DYNAM RENAME MYDD(OLD1,NEW1,R OLD2,NEW2)
DYNAM REN MYDD(OLD1,NEW1 OLD2,NEW2) REPL
```

The SUBMIT Subcommand

The DYNAM SUBMIT command submits jobs to MVS.

The ddname for DYNAM SUBMIT is normally allocated using a DYNAM ALLOC command. On HFS-deployed servers, if a DYNAM ALLOC is not found for the ddname, DYNAM SUBMIT will check whether a FILEDEF was issued for the ddname.

```
DYNAM SUBMIT ddname [MEMBER(members)]
```

where:

SUBMIT

Can be abbreviated as SUB.

ddname

Is the dsname or ddname of the input data set(s) containing JCL to be submitted. The ddname specifies a concatenation of data sets.

MEMBER

May be omitted if the members are specified in parentheses.

members

May be a single member name or a list of members. When a member list is submitted, the resulting job stream is the concatenation of the members. If the members are enclosed in parentheses, blanks before the left parenthesis can be omitted.

Examples:


```
DYNAM SUBMIT MYDD MEMBER ASM,PROG,LKED
DYNAM SUB MYDD(ASM,PROG,LKED)
DYNAM SUB MYID.DATA.LIB(CREATE LOAD)
DYNAM SUBMIT MYFILE
```

The COMPRESS Subcommand

The DYNAM COMPRESS command compresses the partitioned data sets (PDS).

```
DYNAM COMPRESS dname [dname]...
```

where:

COMPRESS

Can be abbreviated as COMP.

dname

Is the dsname or ddname of a PDS to be compressed. The ISPF-like lock is obtained.

If the dsname is specified, it is allocated as OLD. If the ddname is specified and status is SHR, make sure that another user does not access the PDS during the compress operation.

Examples:

```
DYNAM COMPRESS MYDD
DYNAM COMPRESS MYID.DATA.LIB
DYNAM COMP MYDD MYID.DATA.LIB
```

Comparison of TSO Commands, JCL, and DYNAM

This section shows examples of TSO commands and JCL, compared to the equivalent DYNAM commands.

Example: Allocating an Existing File

TSO:	TSO ALLOC F(EDARPC) DA('MYUSER.EDARPC.DATA') SHR
JCL:	//EDARPC DD DSN=MYUSER.EDARPC.DATA,DISP=SHR
DYNAM:	DYNAM ALLOC FILE EDARPC DA MYUSER.EDARPC.DATA SHR

Example: Creating a New Data Set

TSO:	TSO ALLOC F(EDARPC) DA('MYUSER.EDARPC.DATA') - SPACE(5,3) TRACKS CATALOG DIR(2) - UNIT(SYSDA) USING(NEWDCB) - LRECL(80) RECFM(F B) BLKSIZE(1600)
JCL:	//EDARPC DD DSN=MYUSER.EDARPC.DATA,DISP=(NEW,CATLG), // SPACE=(TRK,(5,3,2)),UNIT=SYSDA, // CB=(LRECL=80,RECFM=FB,BLKSIZE=1600)
DYNAM:	DYNAM ALLOC FILE EDARPC DA MYUSER.EDARPC.DATA - SPACE 5,3 TRACKS CATLG DIR 2UNIT SYSDA - LRECL 80 RECFM FB BLKSIZE 1600

Example: Freeing Files

TSO:	TSO FREE F(EDARPC)
DYNAM:	DYNAM FREE FILE EDARPC

Example: Concatenating Files

TSO:	TSO ALLOC F(EDARPC) DA('MYUSER.EDARPC.DATA' - 'MYUSER.PROGRAMS.DATA') SHR
JCL:	//EDARPC DD DSN=MYUSER.EDARPC.DATA,DISP=SHR // DD DSN=MYUSER.PROGRAMS.DATA,DISP=SHR
DYNAM:	DYNAM ALLOC FILE EDARPC DA MYUSER.EDARPC.DATA SHR DYNAM ALLOC FILE PROGRAMS DA MYUSER.PROGRAMS.DATA SHR DYNAM CONCAT FILE EDARPC PROGRAMS

DDNAME Translation in User-Written Subroutines: MSODDX

On z/OS, you can incorporate an additional routine called MSODDX into a user-written subroutine that needs to access ddnames allocated to a Reporting Server. MSODDX provides ddname translation services that enable external programs to access files under the ddname used by the Server.

Reference: Why are DDNAMEs Translated?

Users often want to allocate files according to their individual needs. Therefore, different allocations for FOCEXEC, MASTER, and other necessary DDNAMEs are likely in each Server task. To keep allocations separate for each user, a scheme called DDNAME translation is used to create unique DDNAMEs for every file belonging to each user. When a user allocates a file, the Server internally generates a unique name to perform the allocation and saves the relationship between the DDNAME the user sees and the DDNAME the Server uses for the actual allocation.

For example, the Server FINDMEM subroutine uses the MSODDX routine as follows. When FINDMEM is passed a DDNAME, a member name, and a return code field, the subroutine opens the specified DDNAME to determine if a member with the specified name exists in the file. FINDMEM calls MSODDX to determine the correct DDNAME to open.

Since the allocation performed within the session for a user is not done under the same DDNAME originally specified by the user, when a user-written subroutine needs to access a file, the program must determine what the translated DDNAME is in order to use the actual name when opening the file. This is accomplished by calling the MSODDX routine to perform the translation for you.

Note that only files allocated locally by individual Server users participate in the DDNAME translation scheme. Those present in the Server start-up JCL (IRUNJCL) are referred to as global DDNAMEs since they are available to all of the server tasks. Global DDNAMEs are not translated and, therefore, do not require the MSODDX routine.

MSODDX Calling Environment and Migration From an MVS Server

MSODDX must be called in AMODE31 from an assembly-language program or any high-level language that is LE-compliant and supports dynamic specification of the ddname to be opened. Some languages that do this are PL/I and C. Languages that require the ddname to be hard-coded at compilation time are not suitable (except for COBOL, where TIBCO provides MSODDX2 ASSEMBLER to zap the obtained ddname into Cobol FD). For further information, consult the documentation for your compiler.

If you are migrating from a legacy MVS Server (SSCTL), you must re-link your user-written subroutine to the current version of MSODDX. In addition, if you were previously calling MSODDX in AMODE24, you must now use AMODE31.

Using MSODDX in a Typical Application

MSODDX is typically used to code a user-written subroutine that is callable under the Server and will open and process a file for you. For such a routine, you must call MSODDX once for each file that you open, immediately before issuing the OPEN command.

Using CALLTYPE 1 (as described in [MSODDX Parameter List](#) on page 552), you pass the generic DDNAME—that is, the DDNAME that the user specified when the file was allocated. The correct DDNAME is returned to you. You can either place this into the DCB (using ASSEMBLER), or into a variable (for a high-level language) or, for COBOL, use MSODDX2 ASSEMBLER to zap the obtained DDNAME into a COBOL FD. For details, see [ASMDDX: Sample Assembler Exit for MSODDX](#) on page 554 or [COBDDX: Sample COBOL Exit for MSODDX](#) on page 556.

You can then issue the OPEN command, and the remainder of your application will proceed normally.

Note that it is not necessary to set CALLTYPE=0 in your application to verify that you are running under the Server because if you call MSODDX when not running under the server, MSODDX merely returns the same DDNAME you have passed it.

MSODDX Parameters

MSODDX uses the standard OS-calling conventions. Register 1 must point to a list of full word addresses, which point to the parameters described below. Register 13 must point to a 72-byte register save area. Register 15 contains the entry point address of MSODDX. Register 14 contains the return address to your program. (This is all done for you in a high-level language call.)

Note: For assembler language callers, the return code is provided as one of the parameters and is not given in Register 15 (R15 is returned unchanged).

Reference: MSODDX Parameter List

The following list describes MSODDX parameters and values. The parameters are labeled to indicate whether a value is expected on input (I) to MSODDX, or will contain a value upon return from MSODDX (O).

CALLTYPE

Type of Parameter: Input (I).

Definition: Is the function to be performed. CALLTYPE is a full word (4-byte integer) containing one of the following values:

- ❑ **1** Is the standard call.

You provide the generic DDNAME (for example, FOCEXEC) in GENDD. The correct DDNAME to use is returned in TRANDD.

The RETCODE is:

0 if the DDNAME is known to the Server, either for this individual user, or as a DDNAME known globally to all Server users.

4 if the DDNAME is not known to the Server (TRANDD will be set equal to GENDD so it is not necessary to check RETCODE before using TRANDD).

WORKAREA is required.

- ❑ **2** Means untranslation.

You provide a translated DDNAME in TRANDD. The generic DDNAME it came from is returned in GENDD. If TRANDD is not found to be a translated ddname, GENDD is set equal to TRANDD.

RETCODE is 0 if TRANDD is a translated DDNAME, 4 otherwise.

WORKAREA is required.

- ❑ **3** Creates a translated DDNAME.

Use only when your user-written subroutine will be performing dynamic allocation. Provide the generic DDNAME you intend to use in GENDD. If GENDD is not already in use, a new TRANDD is returned and RETCODE is 0.

If the DDNAME in GENDD is already in use, TRANDD is set equal to GENDD, and RETCODE is set to 4.

WORKAREA is required.

- ❑ **4** Deletes a translated DDNAME previously created using a CALLTYPE 3 call.

Use when your user-written subroutine will dynamically free the DDNAME it had previously allocated. Provide the generic ddname in GENDD.

If found and removed, RETCODE is 0. RETCODE is 4 if an error was encountered.

WORKAREA is required.

GENDD

Type of Parameter: Input.Output (I/O).

Definition: Is the generic (untranslated) DDNAME. GENDD is an eight-byte character field. Its usage is described under the listed CALLTYPE, varying based on the CALLTYPE value.

TRANDD

Type of Parameter: Input/Output (I/O).

Definition: Is the actual (translated) DDNAME. TRANDD is an eight-byte character field. Its usage is described under the listed CALLTYPE, varying based on the CALLTYPE value.

WORKAREA

Type of Parameter: Input (I).

Definition: Is a 300-byte area used by MSODDX for its work area. You must initialize this area to low values (hex zeroes) prior to the first call you make to MSODDX, and should not alter its contents from then on.

RETCODE

Type of Parameter: Output (O).

Is the return code from MSODDX. RETCODE is a full word (4-byte integer) providing an indication of the results of the call made to MSODDX. The only values ever returned will be zero (0) and four (4). Their meanings are described under CALLTYPE, varying based on the CALLTYPE used.

ASMDDX: Sample Assembler Exit for MSODDX

A sample user exit to invoke MSODDX written in Assembler is provided in the *qualif*.HOME.DATA(ASMDDX) library, where *qualif* is the high-level qualifier for your server production data sets.

The same library contains the following:

- ❑ **MSODDXC.** Object code to provide proper LE-compliant addressability between a user-written assembler code and MSODDX.
- ❑ **GENASDDX.** Sample JCL to compile and link ASMDDX and MSODDXC.

You must edit GENASDDX to compile and link the MSODDX routine into your application.

Syntax:

How to Edit GENASDDX

In the following instructions, replace the question marks (?) with the appropriate data set names and high-level qualifiers.

1. Provide a job card.
2. SET INASMLIB='???????.??????.DATA'

This sets the Assembler source library. (For our sample, this corresponds to *qualif.HOME.DATA*.)

3. SET INASMPGM='??????'

This sets the Assembler source name. (For our sample, ASMDDX.)

4. SET INOBJLIB='??????.HOME.DATA'

This sets the Object library where MSODDXC resides (*qualif.HOME.DATA*).

5. SET OUTLIB='??????.??????.LOAD'

This sets the Output load library.

6. SET OUTPGM=ASMDDX

This sets the Output program name.

7. SET SYSLBLK=3200

This sets the Object data set block size.

8. SET PLANG=EDCPMSGE

This sets the Prelinker message module.

9. SET LIBPRFX='CEE'

This sets the LE library prefix.

10. Provide the entry point <AsmEP> in the LKED step.

For the sample ASMDDX, <AsmEP> is ASMDDX.

11. Save GENASDDX, and submit the job.

Example: Testing the ASMDDX Sample

Edit and run the following procedure:

```
SET EXORDER=PGM/FEX
DYNAM ALLOC FILE CAR
EX ASMDDX
```

The ASMDDX load library must be concatenated to DDNAME STEPLIB (in the IRUNJCL that starts the server). Sample ASMDDX output goes to the JESMSGLG output.

Sample Output

In this example, SYS00018 is the translated DDNAME for CAR in the user session on the server.

```
12.43.31 JOB26648 +DD CAR allocated as SYS00018
```

COBDDX: Sample COBOL Exit for MSODDX

A sample user exit to invoke MSODDX written in COBOL is provided in the *qualif*.HOME.DATA(COBDDX) library, where *qualif* is the high-level qualifier for your server production data sets.

The same library contains the following:

- ❑ **MSODDX2**. An Assembler program to ZAP translated DD into the COBOL FD.
- ❑ **GENCPDDX**. Sample JCL to compile and link COBDDX and MSODDX2.

You must edit GENCPDDX to compile and link the MSODDX routine into your application.

Syntax:

How to Edit GENCPDDX

In the following instructions, replace the question marks (?) with the appropriate data set names and high-level qualifiers.

1. Provide a job card.
2. SET INCOBLIB='???????.??????.DATA'

This sets the COBOL source library. (For our sample, this corresponds to *qualif*.HOME.DATA.)
3. SET INCOBPGM='???????'

This sets the COBOL source name.
4. SET INASMLIB='??????.HOME.DATA'

This sets the Assembler source code library where MSODDX2 resides (*qualif*.HOME.DATA).
5. SET OUTLIB='??????.??????.LOAD'

This sets the Output load library.
6. SET OUTPGM=COBDDX

This sets the Output program name.
7. SET LNGPRFX='IBI1.IGY.V2R2M0'

This sets the COBOL compiler library prefix.

8. SET SYSLBLK=3200

This sets the Object data set block size.

9. SET PLANG=EDCPMSGE

This sets the Prelinker message module.

10. SET LIBPRFX='CEE'

This sets the LE library prefix.

11. Save GENCPDDX, and submit the job.

Example: Testing the COBDDX Sample

Edit and run the following procedure:

```
SET EXORDER=PGM/FEX
-* Alloc local DDNAME TESTDD and Global DDNAME CAR.
DYNAM ALLOC FI TESTDD
-RUN
-WRITE TESTDD Line written to TESTDD before invoking COBDDX
-CLOSE TESTDD
DYNAM ALLOC FI CAR NOALIAS
EX COBDDX
```

The COBDDX load library must be concatenated to DDNAME STEPLIB (in the IRUNJCL that starts the server). Sample COBDDX output goes to the SYSOUT DDNAME.

Sample Output

In this example, SYS00003 is the translated DDNAME for TESTDD in the user session on the server.

```
Sample output
TESTING CALL TYPE 0 (WHETHER OR NOT UNDER MSO)
WE ARE RUNNING UNDER MSO (EDA).
RETURN CODE ==> 00000000
TESTING CALL TYPE 1 - STANDARD TRANSLATION CALL
FOR A GLOBAL DDNAME - EXPECT SAME DD AS OUTPUT
GENERIC DDNAME ==> CAR
TRANSLATED DDNAME ==> CAR
RETURN CODE ==> 00000000
TESTING CALL TYPE 1 - STANDARD TRANSLATION CALL
FOR A USER DDNAME - EXPECT TRANSLATION
GENERIC DDNAME ==> TESTDD
TRANSLATED DDNAME ==> SYS00003
RETURN CODE ==> 00000000
```

```
TESTING CALL TYPE 2 - UNTRANSLATION CALL
EXPECT UNTRANSLATION OF DDS USED PREVIOUSLY
GENERIC DDNAME ==> TESTDD
TRANSLATED DDNAME ==> SYS00003
RETURN CODE ==> 00000000
TESTING MSODDX2. ZAPPING DDNAME TO COBOL FD
EXPECT A LINE TO BE READ FROM TESTDD FILE
TRANSLATED DDNAME ==> SYS00003
LINE READ FROM TRANDD ==> Line written to TESTDD before invoking COBDDX
TESTING CALL TYPE 1 - STANDARD TRANSLATION CALL
FOR UNEXISTING DDNAME - EXPECT RETCODE 4
GENERIC DDNAME ==> XXXXXXXX
TRANSLATED DDNAME ==> XXXXXXXX
RETURN CODE ==> 00000004

TESTING CALL TYPE 2 - UNTRANSLATION CALL
FOR GLOBAL DDNAME - EXPECT SAME DD IN GENDD
GENERIC DDNAME ==> CAR
TRANSLATED DDNAME ==> CAR
RETURN CODE ==> 00000004
TESTING CALL TYPE 2 - UNTRANSLATION CALL
FOR UNEXISTING DDNAME - EXPECT RETCODE 4
GENERIC DDNAME ==> XXXXXXXX
TRANSLATED DDNAME ==> XXXXXXXX
RETURN CODE ==> 00000004
```

FILEDEF Command Under UNIX, Windows, IBM i, and z/OS

A logical name (or ddname) is a shorthand name that points to the physical file name as the operating system actually knows the file. Logical names simplify code by allowing short names to be used in place of the longer physical file name.

The FILEDEF command assigns a logical name (or ddname) to a physical file name and specifies file attributes. FILEDEF assignments are in effect for the duration of a connection (except when a server is running in Pool Mode). They are released when the connection to the server is closed or a FILEDEF CLEAR is issued.

FILEDEF with the device name PRINTER is used in conjunction with OFFLINE requests to control printing attributes (such as the printer assignment and number of copies), and is further detailed in *How to Use the FILEDEF Command in UNIX, Windows, IBM i, and z/OS* on page 558.

Syntax: How to Use the FILEDEF Command in UNIX, Windows, IBM i, and z/OS

```
FILEDEF ddname devicetype fileid [[[LRECL n] [RECFM fm] [APPEND]]
```

or

```
FILEDEF ddname DISK app/[appa...]/physfile.ftm
```

```
FILEDEF ddname CLEAR
```

where:

ddname

Is the logical name. It may contain 1 to 8 alphanumeric characters.

devicetype

Identifies the type of device with which to interact. Specifies DISK for a file that resides on disk. Other device types are PRINTER, TRMIN, MEMORY, and TRMOUT, which have special meanings and options. For more information, see [Other FILEDEF Features](#) on page 560.

CLEAR

Clears the specified *ddname*.

app[*appa...*]

Is an application name (APP form of FILEDEF). It can be a nested application name.

physfile.ftm

Is a physical file located in the application (APP form of FILEDEF).

fileid

Is the physical full path name of the file as it is known on the particular operating system, using the native style of that operating system (non APP form of FILEDEF).

For instance, for Windows:

```
c:\mydir\myfile.dat
\\mymachine\mydir\myfile.dat
```

For UNIX, z/OS, and IBM i IFS:

```
/home/myhome/mydir/mtfile.dat
```

For IBM i QSYS

```
QSYS:MYLIB/MYFILE
QSYS:MYLIB/MYFILE(MYMEMB)
```

On UNIX, z/OS, and IBM i IFS file systems, where file names are case-sensitive and lowercase file names are the norm and are expected for the product, appropriate lowercase names should be used when coding the file name. The use of relative file names for server use is not supported, as they are easily confused on some platforms as APP names. The use of an APP MAP command and APP-based naming will let you simulate relative names, yet also be portable to other platforms by simply adjusting the MAP value.

To support physical directory (or file) names with embedded blanks (common on Windows, but also possible on some other operating systems), the complete fileid must be enclosed in single quotation marks. In most instances, a physical operating system name has a unique characteristic that can be used to detect if a physical name is being used so it will not be confused with an APP name. However, IBM i native library names are not always unique enough, so they require an explicit prefix (QSYS:) to ensure proper access.

Note: APP usage can handle any number of directories below APPROOT when nested applications are on (internal default or when the edaserve.cfg file has *nested_app=y*). APP usage is limited to the number of directory levels allowed by the *nested_app* configuration of a server. Any other usage is not permitted.

LRECL *n*

Specifies the record length, *n*, in bytes. This parameter is optional. If you omit it, the default is 80 bytes. Note that the left parenthesis preceding the optional parameters is required.

RECFM *fm*

Describes the record format. Specifies F for fixed format, V for variable format. This parameter is optional. If you omit it, the default is fixed format. Note that the left parenthesis preceding the optional parameters is required.

APPEND

Enables you to open the specified file and add new material at the end of the file. This parameter is optional. If you omit it and the specified file exists, it will be overwritten. Note that the left parenthesis preceding the optional parameters is required.

Note that FOCUS data sources (files with the .foc extension) that do not conform to the default naming conventions are identified using the USE command, not FILEDEF.

Other FILEDEF Features

PRINTER as a device type is used to change the default output file for the OFFLINE print file or set output destinations. For more information, see [Sending Stored Procedure Report Output to Server-Side Printers](#) on page 459.

FILEDEF TRMIN TERM LOWER is used to change the uppercasing behavior of an interactive session (edastart -t) into case-sensitive mode. FILEDEF TRMIN TERM UP is used to restore default behavior. Interactive session mode is typically used for testing and is not considered a production feature for general use.

FILEDEF TRMOUT DISK *fileid* is used to capture session output into a file and is only valid during an interactive session (edastart -t).

FILEDEF TRMOUT TERM is used to restore default behavior after a change to DISK. Interactive session mode is typically used for testing and is not considered a production feature for general use.

FILEDEF *name* MEMORY *file* is used to cache a result in memory. The file must not have any pathing and is typically used for temporary data. The result is gone when the agent completes or CLEAR is issued.

For information about FILEDEF OFFLINE PRINTER, see [Sending Stored Procedure Report Output to Server-Side Printers](#) on page 459.

While command input lines (interactively or in a FOCEXEC) are effectively unlimited, this was not always the case. A feature of FILEDEF to deal with unusually long directory/file names and this former limitation was the ability to split the FILEDEF command across several lines. This was done by using a dash (-) after a FILEDEF token to indicate continuation on the next line. While this feature is not needed anymore, it remains a feature for upward compatibility purposes and is documented here as a deprecated feature.

OFFLINE Printing

Server side printing of formatted reports is accomplished using the OFFLINE command, which sets up and issues a default OFFLINE FILEDEF (or DYNAM ALLOC on z/OS PDS deployment) to receive the formatted outputs after an OFFLINE CLOSE is issued.

The following example on UNIX (and Linux) creates an offline.ftm file and calls lp, the default print command for UNIX (and Linux). On other platforms an equivalent print command is issued:

```
OFFLINE
TABLE FILE CAR
PRINT COUNTRY
END
OFFLINE CLOSE
```

There may be one or more outputs buffered to the same output file for printing, but they are not released to the file until an OFFLINE CLOSE is issued. If a system level variable for FOCPRINT is available at OFFLINE CLOSE time, it will be used to attempt printing of the actual file.

The FILEDEF (or DYNAM ALLOC on z/OS PDS Deployment) OFFLINE PRINTER *device* feature allows the additional specification of output destinations and, in some cases, additional operating system print command switches for features such as multiple copies. In this way, a request can be customized to the print requirements of a particular site.

Prior behavior was that if the operating system variable FOCPRINT was declared (with an operating system command and a \$1), it would be called to take an action on the file name which also replaced the \$1 in the string.

The enhanced FILEDEF command is:

```
FILEDEF OFFLINE PRINTER [filename] [ ( PRINTER printername )
```

The printer name option is used with the standard print feature of a given platform. If no printer name is declared or is set to blank, then the offline file is created, but the print feature is not called. Since printing is platform specific, each platform is described here individually, however, there are still some common needs that are sometimes best resolved by creating a layer that can act as a proxy between the server and the print system.

UNIX and z/OS HFS Deployment

The printer name is dropped in as the "-d" switch value in the "lp -c -d" command. The "-c" switch is used to avoid over-writing of the offline file before actual printing has occurred. If additional lp switches are desired (like multiple copies with the -n switch), they may be stacked into the name by enclosing the string in single quotation marks:

```
FILEDEF OFFLINE PRINTER ( PRINTER '29d1 -n 2'
```

If a site uses lpr instead of lp, then an lp shell script can be created in the \$PATH before the standard lp command and can act as a proxy to call lpr instead. The lp script could be as simple as "/usr/bin/lpr \$*" to redirect lp to lpr.

Note: On z/OS HFS Deployment, the file is spooled to the system and actual disposition will depend on the configuration of the printer spool.

z/OS PDS Deployment

A DYNAM ALLOC command for OFFLINE is used to specify spooling attributes.

Windows

The printer name is a shared printer name and is used to set up (and later drop a NET USE for the LPT1 device to a shared name (for example, \\nodename\myprinter), which is then used in a PRINT /D:LPT1 command to print the actual file. As such, additional switch options cannot be done and use of a PRINT.BAT as a proxy is the only method for further manipulation of the output.

IBM i

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

The output is always spooled from the offline file to the print spool using the system QPRINT file (with whatever the standard values are) on the server's library list. If the spool is set to directly print, output will always be routed as directed with no declaration of a printer name using FILEDEF. If output is not automatically routed and a printer name (using FILEDEF) is declared, then a CHGSPLFA command will be issued with the printer name as the OUTQ() value to direct the spool file to a destination. If additional CHGSPLFA parameters are desired (such as multiple copies with the COPIES() parameter), they may be stacked into the name by enclosing the string in single quotation marks:

```
FILEDEF OFFLINE PRINTER ( PRINTER '29d1 COPIES(2)'
```

A user can also set a specific IBM i QPRINT file to use (which may have attributes such as number of copies set within it) for spooled OFFLINE print files by either placing an appropriate QPRINT file (explicitly named QPRINT) on the server library path or issuing a SET command to use an explicit QPRINT file (of any name). Additionally, you may set a spool file SPLUSRDTA attribute (10-character limited string) to assist in identifying outputs. When set, the print spool job will pick up the attribute information and be displayed by the native operating system tools like DSPLSPLF and WRKDSPLF.

The commands are:

```
SET I5QPRINT = [{library}]{qprintfile}
```

and

```
SET I5SPLUSRDTA = string
```

The internal default for I5QPRINT is QPRINT and may be set to any valid CRTPRTF-created QPRINT file. The QPRINT file and its attributes must be appropriate for the type of file that the OFFLINE command produces (paged plain text) and the capabilities of the actual print device. If no library is supplied, the *LIBL value of the job will be used to locate the file.

Setting I5QPRINT or I5SPLUSRDTA with no value to the right of the equal sign has no effect and leaves the current value unchanged.

Use SET I5SPLUSRDTA = ' ' to set I5SPLUSRDTA explicitly back to no value (blank). Also use single quotation marks for values with embedded spaces.

An I5QPRINT value of blank is illegal for the i5 spool system, so a value of ' ' defaults to QPRINT.

Note that the CRTPRTF command for creating printer attribute files does not create a default QPRINT that matches the IBM-delivered standard default that is in QGPL. To create a QPRINT with standard IBM-delivered default values use:

```
CRTPRTF FILE(*CURLIB/QPRINT) RPLUNPRT(*NO) CHRID(*CHRIDCTL)
```

Then make any site-specific changes and change the owner attribute rights from *ALL to *CHANGE (to prevent over-writing) with:

```
EDTOBJAUT OBJ(*CURLIB/QPRINT) OBJTYPE(*FILE)
```

The use of an lp script on IBM i as a proxy is not effective because lp is not used. The use of alternate QPRINT files is the closest equivalent to an lp proxy.

OFFLINE to DISK Versus PRINTER

The "(PRINTER *printername*" feature is only valid when the FILEDEF device is PRINTER. There is, however, a difference between the use of DISK versus PRINTER as a device in a FILEDEF for OFFLINE. When the device DISK is used, page breaks are represented by a 1 in the first column of a given line where a page break is to occur. This is the FORTRAN Carriage Control method of page control and is a vestige of the product's original mainframe roots. When the disk device is PRINTER, the more modern, Control L (^L / Decimal 12 / Hex 0C) form feed method is used.

Other Printing Information

Very often, sophisticated laser based printers are "hung" off networks and communicated with various print protocols. While these printers may come from many manufacturers, a very common (but not standard) attribute of these types of printers is automatic sensing between a plain clear text file being sent to the printer and a postscript file that contains printer attribute commands as well as the text to print.

OFFLINE files are plain text (vs. HOLD FORMAT PS which do not get spooled using OFFLINE, and it is up to the user to direct them to a printer). Very often, these sophisticated printers can be set up or used improperly, causing a printer to think a plain text file is postscript when it is not, and yielding a page with a postscript error message. This has only been seen so far when printing from Windows, but is in theory possible from any platform.

This problem is not considered to be a server issue because the software is not directly manipulating these printers and uses standard commands supplied by the OS vendor for printing. Generally, this problem can also be reproduced using standard print tools stand-alone from the server environment. A systems administrator for a printer exhibiting an issue like this should be able to track down why this happens in any given environment and take corrective action.

An improper spooling issue may typically also be corrected by creating a proxy script to inject a leading character into the output that resets the printer, so the remaining output properly prints. Generally, this is a control D, but may vary by printer make and model and, as such, specific implementation of such a proxy script is left to the customer since specific needs may vary greatly by site and the various printer models that are available.

Stress Testing a Procedure

Stress testing enables you to simulate request execution in order to ensure that the server will retrieve data satisfactorily under specified conditions.

The mechanism behind stress testing is a record and playback feature in which the exact sequence of user actions applied through a browser are recorded and then reproduced (played back), simulating a single-user or multiple users under the same or different conditions. You can use the generated files and the recorded sequences, known as HTI scripts, along with server traces, to test new configurations, and to diagnose problems.

From the Reporting Server browser interface Playback of HTI Scripts pane, you can initiate test runs under the following separate and combined conditions:

- Requests run simultaneously by a number of users (expressed as threads). You can retest using a variety of scenarios. This feature enables you to determine in advance the system's ability to grow (that is, its scalability).
- Requests executed at varying intervals. This feature enables you to determine in advance the ability of the system to adjust to increases in workload.

In addition to these *basic* tests, you can specify a number of advanced conditions to refine the diagnostic process. For example, stress testing generally provides information about the retrieval performance of the Reporting Server going against one or more data sources (also called back-end processing). If, however, you are working in a client environment like WebFOCUS, you may wish to test data retrieval performance beginning at a user browser, going through a Web server (front-end processing), to the Reporting Server to the data source. You can request that the test be redirected to the Web server to track retrieval over the longer path. Then, by comparing back-end test results with front-end test results, you can more accurately diagnose the origin of a performance problem based on any other conditions you specify.

For a full list of the basic and advanced stress test conditions you can apply, see [Parameters for Playback of HTI Script Files](#) on page 787.

Recording User Actions Into a Script Using Run Stress

The server recording feature can record the exact sequence of user actions applied through a browser into a script. The script can then reproduce that sequence (play it back), simulating single or multiple users under varying conditions to produce execution statistics.

The Run Stress feature uses an existing procedure to produce the script. It also changes the browser session to the Diagnostics Scalability Playback Start/Stop page for the script.

The files used by playback and the recorded sequences are known as HTI scripts and are saved in the scalability directory. The scripts can be used along with server traces for problem diagnostics, analysis, and testing.

You can go back to the Diagnostics Scalability Playback page at any time off the main menu and create a new script by recording a session activity (top menu bar) or an edit session (top menu bar) plus edit or re-run an existing script.

Procedure: How to Run a Stress Test

You can run a stress test on a procedure from the Reporting Server browser interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a procedure, point to *Run Advanced*, and click *Run Stress*.

The Playback of HTI Script Files dialog box opens, as shown in the following image.

The image shows a dialog box titled "Playback of HTI Script Files" with a "Basic" tab selected. The "Basic" tab contains the following fields and options:

- Script Directory:** A text box containing "ibisamp".
- Script Name:** A text box containing "mrsamp01".
- Number Of Threads:** A text box containing "1".
- Interval Parameter:** A text box containing "0.5".
- Keep alive for:** A text box containing a vertical bar "|", followed by the text "seconds".
- Delay:** A text box, followed by the text "seconds between requests in keep-alive mode".
- Test Description:** Three stacked text boxes, all of which are empty.

Below the "Basic" tab is an "Advanced" tab, which is currently collapsed. At the bottom of the dialog box is a "Start" button.

4. Enter basic test conditions: Number of Threads, Interval Parameter, and a Test Description. Click *Advanced* if you wish to specify additional conditions.
5. Click *Start*. Results are displayed in a separate window. For an illustration, see [Sample Stress Test Log](#) on page 568.

The script is listed under the tested procedure in the directory C:\ibi\scale.

6. Open the Performance Reports folder on the Playback navigation pane and choose:
 - Basic Report*
 - Extended Report*
 - Graph Report*

Select *View* to view the report or *Open* to see the report or the Metadata, Query, and Output panes for the report.

Example: Sample Stress Test Log

---- Started at 13:39:09 ----

Received: thread=01 request=0001; timing: resp=0.052 sec, transf=0.000 sec,
start=03/27/2018 13:39:09.795; dbmstime=0.000 sec, servertime=0.052 sec

---- Finished at 13:39:09 ----

Total Execution Time: 0.115 sec
 Total Number of Threads: 1
 Interval Parameters: 500,-1,-1,-1
 Minimum Execution Time of Thread: 0.109 sec
 Maximum Execution Time of Thread: 0.109 sec
 Average Execution Time of Thread: 0.109 sec

Total Number of Requests: 1
 Average Server Response: 0.052 sec of total 1 requests
 Average Data Transfer Time: 0.000 sec of total 0 requests
 Average Processing Time: 0.052 sec of total 1 requests
 Standard Processing Deviation: 0.000 sec
 Minimum Processing Time: 0.052 sec
 Maximum Processing Time: 0.052 sec
 Average DBMS Time: 0.000 sec of total 1 requests
 Minimum DBMS Time: 0.000 sec
 Maximum DBMS Time: 0.000 sec
 Average Server Time: 0.052 sec of total 1 requests
 Minimum Server Time: 0.052 sec
 Maximum Server Time: 0.052 sec

terminating main thread

Sample Performance Reports

Basic Report

Basic Stress Test Statistics


Test Name	Test Description 1	Begin Time	End Time	Total Time	Arrival Rate per minute	Average Response Time (sec)	Average DBMS Time	Maximum Server Time (sec)
fbisamp/flow2xml		2017/04/26 17:12:43	2017/04/26 17:12:43	.123	120.000	.072	.000	.072
fbisamp/flow2xml		2017/04/26 17:13:37	2017/04/26 17:13:48	10.891	120.000	.113	.000	3.875
fbisamp/experson		2017/04/26 17:15:36	2017/04/26 17:15:46	9.931	120.000	.045	.000	.075

Extended Report

Extended Stress Test Statistics

Test Name	Test Description 1	Test Description 2	Test Description 3	Test Description 4	Begin Time	End Time	Total Time	Arrival Rate per minute	Total Number of Requests	Total Number of Threads	Interval Parameters	Average Response Time (sec)
fbisamp/flow2xml					2017/04/26 17:12:43	2017/04/26 17:12:43	.123	120.000	1	1	500;-1;-1;-1	.072
fbisamp/flow2xml					2017/04/26 17:13:37	2017/04/26 17:13:48	10.891	120.000	264	3	500;-1;-1;-1	.113
fbisamp/experson					2017/04/26 17:15:36	2017/04/26 17:15:46	9.931	120.000	211	1	500;-1;-1;-1	.045

Graph Report



Running Impact Analysis Reports

An Impact Analysis report enables you to determine the impact of changing a synonym, synonym column, procedure, or connection. The report shows you how and where the object is used. For example, running the report on a procedure tells you from which applications and procedures the original procedure is run. Running the report on a synonym tells you which tables or other data sources the synonym describes, and in which procedures it is referenced. This enables you to determine how changing the synonym, column, procedure, or connection could affect the rest of your application.

You can run Impact Analysis reports from the Reporting Server browser interface and from the Data Migrator desktop interface. Reports also provide drill-down links so that you can, for example, run an Impact Analysis report on a synonym, see a procedure listed in the report, and click the procedure to edit it.

In addition to the reports you can run for an individual procedure, synonym, connection, or column, you can select the following standard Impact Analysis reports for an Application folder in the Reporting Server browser interface or the Data Migrator desktop interface:

- The **Synonyms by Procedure** report lists your procedures and the synonyms that they use, sorted by procedure. The report includes whether the synonym is used as a source or a target, the adapter type, the table name and description, load type and prior to load options, and directory locations.

- The **Procedures by Synonym** report lists your synonyms and the procedures they are used in, sorted by synonym. The report includes whether the synonym is used as a source or a target, the adapter type, the table name and description, load type and prior to load options, and directory locations.
- The **Columns by Procedure** report lists the columns used in your procedures and their synonyms sorted by data flow. The report includes how the column is used, as well as directory locations.
- The **Procedures by Column** report lists your procedures and their synonyms sorted by column. The report includes how the column is used, as well as directory locations.
- The **Synonyms by Procedure for Enterprise** report includes the same information as a Synonyms by Procedure report, but for the server and all its subservers in a hub/subserver environment. This report does not appear if the server you are reporting on does not have any remote servers configured.
- The **Procedures by Synonym for Enterprise** report includes the same information as a Procedure by Synonym report, but for the server and all its subservers in a hub/subserver environment. This report does not appear if the server you are reporting on does not have any remote server configured.
- The **Flow Report** report has two options, Summary and Detailed, which report on existing flows.

Note: Enterprise reports appear only if you have subservers.

***Procedure:* How to Run an Impact Analysis Report for an Application**

You can run an Impact Analysis Report from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Right-click an application folder, point to *Impact Analysis*, and click one of the following options:
 - Synonyms by Procedure*
 - Procedures by Synonym*
 - Columns by Procedure*
 - Procedures by Column*

□ Flow Report

The Impact Analysis report opens in the right pane.

Procedure: How to Run an Impact Analysis Report for an Adapter Connection

You can run an Impact Analysis Report from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Get Data* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an adapter folder, if necessary, or click *Show Connections*.
3. Right-click a connection, and select *Impact Analysis*.

A report similar to the following image opens in the right pane. The report provides links to the files and lists such information as the connection names, type of file, usage, and application name.

Found In Application	Found In	Found In Type	Connection	Usage	Adapter	Line	Procedure Description
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_currency	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Currency Dimension
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_customer	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Customer Dimension
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_discount	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Discount Dimension
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_product	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Product Dimension
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_store	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Store Dimension
ibisamp/se/dimensions	pivot_demo_wrd_wf_retail_time	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Time/Date Dimension
ibisamp/se/facts	pivot_demo_wrd_wf_retail_sales	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Sales Fact
retail8203/bv_namespace_off/advanced	wf_retail_product_base	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	2	Product BV
retail8203/bv_namespace_off/advanced	wf_retail_sales_base	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	2	Sales Fact
retail8203/bv_namespace_on/advanced	wf_retail_product_base	Business View	CON02	CONNECTION-	MS SQL Server ODBC	2	Product BV
retail8203/bv_namespace_on/advanced	wf_retail_sales_base	Business View	CON02	CONNECTION-	MS SQL Server ODBC	2	Sales Fact
retail8203/dimensions	wf_retail_age	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Age Dimension
retail8203/dimensions	wf_retail_currency	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Currency Dimension
retail8203/dimensions	wf_retail_customer	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Customer Dimension
retail8203/dimensions	wf_retail_discount	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Discount Dimension
retail8203/dimensions	wf_retail_education	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Education Dimension
retail8203/dimensions	wf_retail_employee	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Employee Dimension
retail8203/dimensions	wf_retail_geography	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Geography Dimension
retail8203/dimensions	wf_retail_income	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Income Dimension
retail8203/dimensions	wf_retail_industry	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Industry Dimension
retail8203/dimensions	wf_retail_labor_rate	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Labor Rate Dimension
retail8203/dimensions	wf_retail_marital_status	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Marital Status Dimension
retail8203/dimensions	wf_retail_occupation	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Occupation Dimension
retail8203/dimensions	wf_retail_product	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Product Dimension
retail8203/dimensions	wf_retail_shipping_options	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Shipping Options Dimension
retail8203/dimensions	wf_retail_store	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Store Dimension
retail8203/dimensions	wf_retail_time	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Time/Date Dimension
retail8203/dimensions	wf_retail_time_lite	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Time/Date Dimension
retail8203/dimensions	wf_retail_vendor	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Vendor Dimension
retail8203/facts	pivot_demo_wrd_wf_retail_shipments_2	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Shipments Fact
retail8203/facts	wf_retail_labor	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Labor Fact
retail8203/facts	wf_retail_sales	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Sales Fact
retail8203/facts	wf_retail_shipments	Synonym	CON02	CONNECTION-	MS SQL Server ODBC	1	Shipments Fact

Procedure: How to Run an Impact Analysis Report for a Procedure

You can run an Impact Analysis Report from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.

3. Right-click a procedure, and click *Impact Analysis*.

The Impact Analysis report opens in the right pane.

***Procedure:* How to Run an Impact Analysis Report for a Synonym**

You can run an Impact Analysis Report from the Reporting Server browser interface or the Data Migrator desktop interface.

1. Go to the Reporting Server browser interface *Applications* page, or from the Data Migrator desktop interface, expand the Server node folder.
2. Expand an application folder.
3. Right-click a synonym, and click *Impact Analysis*.

The Impact Analysis report opens in the right pane.



Chapter 8

Workspace Manager

The Reporting Server browser interface provides a variety of tools and techniques that you can use to run and monitor your server. The Workspace Manager is the controlling process that controls and oversees the various processes of the server.

In this chapter:

- [Configuring Workspace Manager](#)
 - [Notification](#)
 - [Workspace Configuration Settings](#)
 - [Monitoring TIBCO WebFOCUS Reporting Server Activity](#)
 - [Editing Configuration Files](#)
 - [Configuring Java Services](#)
 - [Cluster Manager](#)
-

Configuring Workspace Manager

The Workspace Manager is the component of the Reporting Server that is responsible for managing all server administrative tasks. These tasks, generally performed by Server Administrators, include monitoring server activity, configuring and adjusting the server configuration profile, enabling and creating services, defining deferred execution characteristics, and enabling and disabling email alerts.

The Server Administrator responsible for the installation, configuration, security, and maintenance of the server uses the Reporting Server browser interface to manage and configure the Workspace Manager in order to keep the server available to clients and running at peak efficiency. To use the Reporting Server browser interface, open Internet Explorer and navigate to the HTTP port on the host machine where the Server is running, as in `http://host:http_port`.

Your Server must be running in order to use Reporting Server browser interface to access the Workspace Manager.

Tip: This chapter references a number of Workspace related keywords. Within the help system, you can access detailed information about these keywords by clicking the keyword links in this document. You can also click the question mark (?) icon next to parameters on the Workspace configuration panes.

Workspace Manager Configuration Privileges

Access to the administrative features of the Reporting Server browser interface can be restricted by defining a list of users with admin privileges and storing the list in a release-independent file called `admin.cfg`. This file is located, by default, in `.../ibi/profiles`.

The list of users and user roles defined in the `admin.cfg` file defines a list of administrators and users that can be used for authorization and/or authentication according to established security. Administrators are responsible for installing, configuring, and maintaining the Server with varying degrees of responsibility, depending on their administration levels (SRV, APP, OPR). At least one administrator must be defined in the list, although most sites identify other persons to act as backups.

For example, a Server Administrator (SRV) can perform all the administrative tasks available through Reporting Server browser interface operations. If more than one Server Administrator is defined, the first valid member on the list is used to impersonate FOCUS Database Server (FDS) and other special services. An Application Administrator (APP) is limited to the administrative tasks that do not require changing configuration or restarting the server. Both Server Administrators and Application Administrators can edit user profiles in the `user.prf` file. However, Server Administrators can edit all `user.prf` files, while Application Administrators can only edit their own profiles.

Any IDs (beyond the original ID used to configure the server) that are used for server or application administration require read/write privileges to the respective locations that the IDs are expected to manage. To assign these privileges, you must establish group rights for the locations at the operating system level.

Configuration actions depend on privileges. For example, Resource Management procedures are available to users with the RARGP privilege.

The Server Actions menu on the Workspace Manager ribbon has options to stop the server, restart the server, and quiesce the server.

Stop Server

A server administrator can stop the server. Existing connections are terminated.

Restart Server

A server administrator can restart the server. Existing connections are terminated, and new connections will be rejected until the server restarts.

Quiesce/Enable Server

A server administrator can change the server state to quiesce to disable new connections. Existing connections are not terminated, but will finish processing and disconnect, while new connections will be rejected. A server administrator can still connect to the server. This is done to gradually move the server to the state where maintenance can be performed safely or to collect diagnostic information. It can also be used as the first step before stopping the server without terminating any connections.

A server administrator can enter a custom message that will be displayed for a new user connection when the server is in Quiesce mode.

Procedure: How to Quiesce the Server and Set a Custom Message

Only a server administrator can change the server state. To Quiesce the server, and define a custom message that will be delivered from the server:

1. From the Server Actions menu on the ribbon, click *Quiesce/Enable*.

The Quiesce Connections page opens.

2. If you want to deliver a message for new user connections, enter it in the `quiesce_msg` field.

Server is currently accepting all connections.

? `quiesce_msg`

Connection to server is temporarily disabled by administrator

Submit

Cancel

3. Click *Submit*.

You will be asked if you want to disable new connections.

4. Click *OK*.

Note: To restore the server to normal operation, click *Quiesce/Enable*, and select *Enable New Connections*.

Configuring an Agent Service

A server configuration requires at least one agent service with the name `DEFAULT`, defined by a `SERVICE` block. An agent service is an entity used to define the parameters for a group of data access agents, so that a configuration can manage different groups of data access agents for different purposes. Each data access agent runs for a specific Data Service, and each service may have different values for the settings defined on the services configuration panes. These settings include:

- The maximum number of data access agents and the number of agents prestarted at server startup for a service, as defined by *maximum* and *number ready*.
- The lifetime of the agents of a server, which can be limited through *idle agent limit*, *CPU time limit*, and/or *memory limit*.
- Incoming connections for which there is no available data access agent. These can be put in a queue for the service (configured using *maximum_q* and *queue limit*), and after they are connected, their idle time can be limited using *idle session limit*.
- The *deployment mode* of a service, which defines how data access agents are assigned to connections.

Procedure: How to View or Edit Data Service Properties

The service agents are available from the *Workspace* resources tree.

1. Go to the *Workspace* page.
2. In the resources pane, open *Data Services*.
3. Right-click a Data Service and select *Properties*.

The properties page for the particular Data Service opens.

Reference: Data Service Properties

The server includes four predefined Data Services:

- `DEFAULT`
- `WC_DEFAULT`
- `SCHED_DEFAULT`

❑ DFM_DEFAULT

The following image shows the parameters for the DEFAULT service.

Changing values with an asterisk (*) requires restarting the server

? service	DEFAULT
? maximum	<input type="text" value="40"/>
? number_ready	<input type="text" value="10"/>
? deployment *	<input type="text" value="private"/>
? Queuing *	<input type="text" value="off"/>
? idle_session_limit	<input type="text" value="-1"/>
? idle_agent_limit	<input type="text" value="120"/>
? profile *	<input type="text"/>
? cpu_limit	<input type="text" value="-1"/>
? memory_limit	<input type="text" value="-1"/>
? disk_limit	<input type="text" value="-1"/>
? connection_limit	<input type="text" value="-1"/>
? max_connections_per_user *	<input type="text" value="-1"/>
? agent_refresh	<input type="text" value="100"/>
? sched_priority	<input type="text" value="0 (normal priority)"/>

Note: Changing values with an asterisk (*) requires restarting the server.

Data Services have the following parameters:

service

Defines the name of the service.

maximum

Defines the maximum number of data access agents that the Workspace Manager will allow to run simultaneously for a specific service.

`number_ready`

Defines the number of data access agents that the Workspace Manager will create at startup for a specific service.

`deployment*`

Controls resource sharing between users who are connecting to the same agent in this service. The values are:

- private deployment.** All users are completely independent of each other and have their unique operating system rights, database connections, and FOCUS language settings.
- connection pooling.** All users are partially independent by having their own unique FOCUS language settings but sharing operating system rights and database connections.

`Queuing*`

Controls whether queuing is on or off.

`maximum_q`

Defines the maximum number of connections that could be queued for a specific service. Only available when queuing is set to on. Set to -1 for an unlimited queue size.

`queue_limit`

Defines the amount of time in seconds a queued connection will wait before being timed out if an agent is still unavailable. A setting of -1 mean unlimited.

`idle_session_limit`

Defines the time limit in seconds that connected agents will wait for client input before they are disconnected. A setting of -1 mean unlimited.

`idle_agent_limit`

Defines the time limit in seconds that disconnected agents in excess of `number_ready` can stay idle before they are killed. A setting of -1 mean unlimited.

`profile*`

Specifies a focexec file that will be executed during agent startup.

`cpu_limit`

Defines the amount of CPU time in seconds an agent is allowed to use before being killed by the Workspace Manager.

`memory_limit`

Defines the maximum amount of memory in kilobytes an agent is allowed to use. If an agent process grows above this limit, it will be killed by the Workspace Manager.

`disk_limit`

Defines the maximum amount of disk space in kilobytes an agent is allowed to use. This would add sizes of all files in the agent edatemp directory, such as FOCSORT, HOLD files and other temporary files created by requests. If an agent process grows above this limit, it will be killed by the Workspace Manager.

`connection_limit`

Defines the time limit in seconds allowed for connection. When it is exceeded, the connection will be terminated, and the agent serving this connection will be stopped. The limit will be triggered after the current edachkup_interval expires.

The edachkup_interval is an internal parameter that defines the time interval for the Workspace Manager checkup cycle. Data service level limit parameters such as connection_limit, idle_agent_limit, cpu_limit, and disk_limit take effect the next time edachkup_interval wakes up (by default, every 60 seconds).

`max_connections_per_user*`

Restricts the number of concurrent connections for the same userid for a specific service. When a user exceeds the maximum number allowed, additional connections requested by that user are rejected, and the server displays the message

```
Connection refused due to the max_connections_per_user (n)
being exceeded
```

where n is the number of allowed connections.

You have the option to queue user connections that are refused by the server, using the queue_max_user_conns property on the Data Service Properties page. That option is available only if the Queuing property for the Data Service is *On*.

`queue_max_user_conns`

Enables queuing of concurrent connections for the same userid for a specific service in excess of max_connections_per_user. This setting allows new connections that would otherwise be rejected to be queued for later processing. Only available when queuing is set to on.

`agent_refresh`

Defines the maximum number of new connections which can be accepted during the life of each agent process. Beyond this limit, additional new connections will be assigned to a fresh agent. Agent processes which have reached the limit will be terminated when the last accepted session disconnects.

`sched_priority`

Defines the scheduling priority of each agent process running under the specified service entry. The scheduling priority of a process defines how the operating system scheduler treats the process after it gains control of the CPU, and should be set according to the relative importance of the work being done by the service. Values range from -20 (highest priority) to 20 (lowest priority).

Reference: How the Reporting Server Maintains the Number of Available Agents

The Reporting Server maintains the number of available agents using the `number_ready` parameter.

- ❑ The `number_ready` parameter will determine how many agents are started when the Reporting Server starts. However, it does not guarantee that the number of active agents will not fall below that value. The number of agents may be lower if agents crash or are stopped unexpectedly. Even if the number of active agents is lower than the `number_ready` parameter, more will be started as needed, up to the number allowed by the Maximum parameter.
- ❑ If there are active agents that are idle, and that number exceeds `number_ready`, agents will be stopped once they exceed the `idle_agent_limit` until the number of active agents equals the `number_ready`.

Configuring Deployment Modes

The *deployment mode* of a service defines how data access agents are assigned to connections:

- ❑ In private deployment, a dedicated application agent is assigned for each connection request. Private deployment retains the behavior of all prior server releases. At connect time, `global`, as well as `user` and `service level` profiles, are executed. At disconnect time, all temporary files are removed and database connections are closed. The privileges of each application agent depend on the security mode of the server.

With security set to a value other than OFF, authentication is processed for every client logging on to the server. With security OFF, user identification and authentication are not required. The effective user ID becomes the connecting user for the duration of the session.

- ❑ In connection_pooled deployment, the global profile and service profiles are executed when an agent is started, and pooled user profiles are executed on each connection. The WebFOCUS context is cleared once the session is established for a new connection, and then the pooled user profile is executed.

With security set to a value other than OFF, all users have the same rights because the effective pooled user is unique, regardless of the connecting user.

Connection pooled deployment provides significant performance advantages and is recommended when a large number of users share the same operating system and DBMS credentials. It cannot be used when each connecting user has specific operating system and DBMS rights, nor can it be used for the service DEFAULT with the server security mode set to DBMS because, in this situation, the profile is disabled for connection authentication.

Procedure: How to Set the Server Deployment Parameter in a Server Profile

To set the server configuration parameter deployment:

1. Go to the *Workspace* page.
2. In the resources pane, right-click the *Data Services* folder, and click *New*.

The New Data Service page opens.

Tip: For an existing service (for example, DEFAULT, WC_DEFAULT, DMC_DEFAULT, DEFAULT_CPOOL), right-click the service and click *Properties* to open the Services pane.

3. In the deployment field, choose *private* or *connection_pooling*.

- ❑ With private deployment, all profiles are executed on connection. On disconnect, all DBMS connections are dropped. This is the default value.

With security set to a value other than OFF, the effective user ID is switched to the connecting user for the duration of the connection.

Skip to step 6.

- ❑ With connection-pooled deployment, the global profile and service profiles are executed when an agent is started, and pooled user profiles are executed on each connection. The WebFOCUS context is cleared once the session is established for a new connection, and then the pooled user profile is executed.

When you select this option, *pooled_user* and *pooled_password* are displayed.

With security set to a value other than OFF, all users have the same rights because the effective user is unique (configured using *pooled_user* and *pooled_password*) regardless of the connecting user.

Continue with steps 4 and 5.

4. Select *pooled user* from the drop-down list. (This list is populated from the *admin.cfg* file described in [Workspace Manager Configuration Privileges](#) on page 574.)

This service level keyword is required for *connection_pooling* deployment with security mode OPSYS. It defines the user ID under which all agents will run. The DBMS user IDs are determined by the connection setting type.

5. Enter the corresponding *password*. Note that pooled password is only required for *connection_pooling* deployment on Windows operating systems.
6. Click *Save and Restart Server*.

Notification

Event Routing and a customized Login Message can be configured from the Notification menu on the Workspace page.

Event Routing allows you to launch procedures or send an email based on different server events (agent or listener crashes, server configuration errors, disk space limitations, and others). All events are classified by type as information, warning, or error. These can be pre-defined or custom-defined events.

The Event Routing page shows the Default and Customized Event Routing.

The Default routings are shown in the following image.

Type	Filter	Description	EDAPRINT	Terminal	Procedure	E-Mail
- Default						
Information	All Information		✓	✓		
Warning	All Warnings		✓	✓		
Error	All Errors		✓	✓		
Information	13023	Server successfully started	✓	✓		
Information	13093	User password has been changed	✓	✓		
Information	13120	Server is requested not to accept any new connections	✓	✓		
Information	13121	Server is requested to accept all connections	✓	✓		
Information	13164	A listener is enabled	✓	✓		
Information	13163	A listener is disabled	✓	✓		
Information	13169	A data service agent has been started	✓	✓		
Information	4276	DBCTL interface has been initialized	✓	✓		
Information	4288	DBCTL interface is stopped	✓	✓		
Warning	13088	Connection is rejected because userid is invalid	✓	✓		
Warning	13089	Connection is rejected because password is invalid	✓	✓		
Warning	13090	Failed to change user password	✓	✓		
Warning	13092	Connection is rejected because account is expired	✓	✓		
Warning	13173	Number of data service agents has reached the maximum allowed	✓	✓		
Warning	40001	No more available agents for required service (maximum number is reached)	✓	✓		
Warning	40002	Not enough disk space in dfm_dir directory to process deferred requests	✓	✓		
Warning	40003	File listener warning	✓	✓		
Warning	40004	CLM failed to query remote server	✓	✓		
Warning	40005	Agent exceeded CPU limit	✓	✓		
Warning	40006	Agent exceeded memory limit	✓	✓		
Warning	40007	Agent exceeded disk space limit	✓	✓		
Warning	40008	Agent exceeded connection time limit	✓	✓		
Warning	36048	Resource Governor canceled a request	✓	✓		
Warning	36049	Resource Governor advised a request	✓	✓		
Warning	36037	Resource Management has unarchived logs	✓	✓		
Information	36162	Resource Management has disabled monitoring	✓	✓		
Information	36163	Resource Management has enabled monitoring	✓	✓		
Error	36029	Resource Management repository validation failure.	✓	✓		
Error	36036	Resource Management Archive job parameter validation failed	✓	✓		
Error	36075	Resource Management Archive job failed, no admin defined	✓	✓		
Error	36076	Resource Management Archive job connection failed	✓	✓		

Procedure: How to Customize an Event Routing

By customizing event routing, you can launch a procedure or send an email based on server events.

For each event, you can also select whether to write the event to the EDAPRINT log or the terminal.

1. Go to the *Workspace* page.
2. Click *Notification*, then *Event Routing*.
3. Right-click an event and click *Customize Event Routing*, or click *New* on the menu bar.

The Customize Event Routing page opens, as shown in the following image.

? Type:

? Filter:
 Select pre-defined filter or type in a new one (Do not use commas in filter)

? EDAPRINT

? Terminal

To

? E-Mail Cc

Bcc

Use semicolon to separate emails, user IDs, procedures.

? Procedure

4. From the Type drop-down menu, select a message type. The choices are *Information*, *Error*, *Warning*, and *Any*.
5. Enter the event criteria in the Filter field, or select a predefined event from the drop-down menu.
6. From the EDAPRINT drop-down menu, select Yes if you want to write the selected event message to the server log.
 In z/OS, the selected server messages can be also written to the z/OS system log by selecting Yes from the JESLOG drop-down menu.
7. From the Terminal drop-down menu, select Yes if you want to write the selected event message to the terminal.
8. Optionally, enter email addresses to send an email when the event occurs.
9. Optionally, enter a procedure to execute in the Procedure field, including the application name, for example, utility/movefiles.
10. Click *Add* for a new event, or *Change* for an existing event.

The event is added to the Customize Event Routing Page.

11. Click *Save and Restart Server*.

Example: Suppressing an EDAPRINT Log Message

This example shows how to suppress the following EDAPRINT log message by customizing an event routing:

```
request by edapsmr to expire foccache session
```

1. Open the Event Routing page.
2. Select *Warning* from the *Type* drop-down list.
3. Enter the following event in the *Filter* text box:


```
request by edapsmr to expire foccache session
```
4. Select *No* from the *EDAPRINT* and *Terminal* drop-down lists.
5. Leave the E-Mail parameters blank, so that an email will not be sent when the event occurs.
6. Click *Add*.

The event is added, but not yet saved.
7. Click *Save and Restart*.

When the server restarts, this custom event will suppress the *request by edapsmr to expire foccache session* message in the *edaprint.log* file.

Procedure: How to Set Workspace Log Properties for Event Routing

The Workspace Log Properties page controls both event routing email and the edaprint log by allowing you to specify a suppression interval for email notifications from event routing, and to include a number of lines from the edaprint log file in the email body. The edaprint log file provides information about the events occurring on the server before a triggered event.

For information on the edaprint log, see [How to Set Workspace Log Properties to Control the Edaprint Log](#) on page 762.

1. Go to the *Workspace* page.
2. Open *Logs and Traces*.
3. Right-click *Workspace Logs* and click *Properties*.

The Workspace Logs Properties page opens, as shown in the following image.

The screenshot shows a web interface with a breadcrumb trail: Data Services Agents > New Data Service > Workspace Logs Properties. There are two main sections: 'Event E-Mail' and 'EDAPRINT'. The 'Event E-Mail' section has two input fields: 'email_suppress_interval' with the value 600 and 'logging_email_edaprint' with the value 20. The 'EDAPRINT' section has four controls: 'edaprint_history' (input field with 5), 'edaprint_max_lines' (input field with 0), 'edaprint_max_days' (input field with 0), and 'edaprint_ipname' (dropdown menu with 'n' selected). At the bottom right, there are 'Cancel' and 'Save' buttons.

4. Enter or select values for the Event Email parameters.

The Event Email parameters are:

`email_suppress_interval`

Suppresses duplicate email notifications if the same event occurs again during this time interval (in seconds). Zero disabled suppression. The default value is 600 (10 minutes).

`logging_email_edaprint`

Defines the maximum number of server messages from edaprint.log are included in the email. The default value is 20.

5. Click Save.

EVENT_TEXT and EVENT_ID

When a procedure is triggered by an event, the ID of the routing event and the text of the event message are automatically available to the procedure as the `EVENT_ID` and `EVENT_TEXT` parameters, respectively. The `EVENT_TEXT` parameter needs to be decrypted using the following user subroutine:

```
B64DECODE('XXXXXXXXXX', 'Ann');
```

Login Message

Server administrators can customize the login and connection result screens with custom messages. The *Login Message* option is available from the *Notification* button on the ribbon.

Procedure: How to Customize the Login Screen With a Message

To customize the login screen:

1. Go to the *Workspace* page.
2. Click *Notification*, then *Login Message*.

The Login Message page opens, as shown in the following image.

The screenshot shows a configuration interface for login messages. It is organized into two main sections: 'Before' and 'After'.
 - The 'Before' section contains a 'HTML Message' text area. Below it, an example is provided: `click me`.
 - The 'After' section contains a 'Role' dropdown menu currently set to 'Basic User', a large empty 'Message' text area, and a 'Display' field with the value '1' and the label 'time(s) to each user'.
 - At the bottom of the form are two buttons: 'Save' and 'Cancel'.

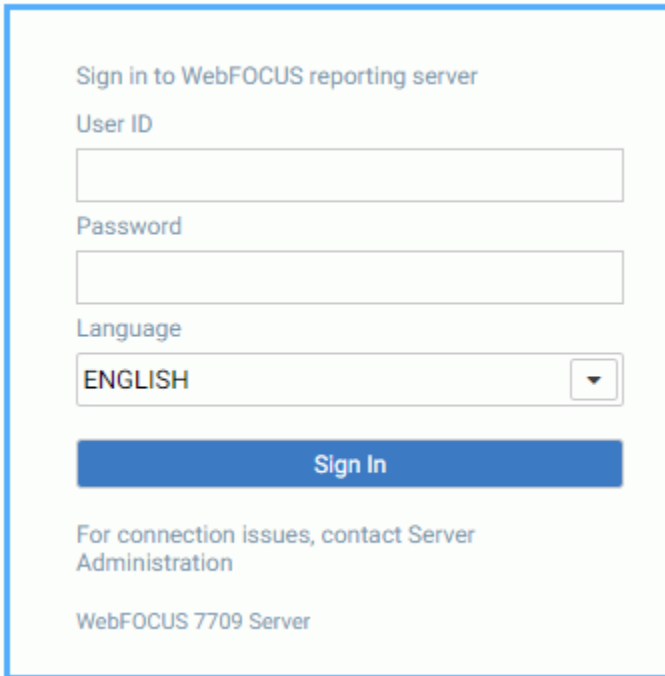
3. Enter the message in the HTML Message field in the Before section. It can have plain text as well as HTML tags.

For example, enter the following text.

[For connection issues, contact Server Administration](#)

4. Click Save.

The message is added to the login page, as shown in the following image.



The image shows a login form for the WebFOCUS reporting server. The form includes fields for User ID, Password, and Language (set to ENGLISH). A blue Sign In button is present. Below the button, a custom message is displayed: "For connection issues, contact Server Administration" and "WebFOCUS 7709 Server".

Procedure: How to Customize the Connection Result Screen With a Message

Server administrators can customize the connection result screen for each user type.


1. Go to the *Workspace* page.
2. Click *Notification* then *Login Message*.
The Login Message page opens.
3. Select a user type from the Role drop-down menu. The options are *Server Administrator*, *Application Administrator*, *Server Operator*, *Basic User*, and *All Users*.
4. Enter the text of the message in the Message field of the After field:

Please sign off by 6pm tonight.

Optionally, you can set the number of times the message will appear by entering a value in the Display field.

5. Click Save.

The message is displayed after a successful connection, as shown in the following image.



Workspace Configuration Settings

You can set a variety of general parameters from the *Workspace* folder in the resources pane. The parameters are available from a right-click menu and from buttons, organized under the following categories:

Notification

- [Notification](#) on page 582.
- [Login Message](#) on page 587.

Settings

- [Reporting Server Web Console Appearance](#) on page 590.
- [Locale Settings](#) on page 591.
- [Displaying Settings for Reporting Server Browser Interface Preferences](#) on page 601
- [TIBCO FOCUS Sets and Info](#) on page 605.
- [Workspace SET](#) on page 617.
- [Configuring Geographic Information](#) on page 627.
- [Migrate Settings](#) on page 658.

These settings are saved in the appropriate configuration files. You can also save them in a procedure and run them by making the procedure a service profile (configured in the service properties) or run them by executing the procedure from a configuration file.

Tip: If you prefer, you can edit these parameters directly in the associated configuration files. To do so, go to the *Workspace* page. On the resources pane, expand *Configuration Files*, and choose the option that opens the file you wish to edit in a text editor. For details, see [Editing Configuration Files](#) on page 701.

You can monitor and troubleshoot server activity from *Workspace* in the resources tree. These functions are available from a right-click menu and from options on the ribbon, organized under the following categories:

Monitor

- Status. See [Monitoring TIBCO WebFOCUS Reporting Server Status](#) on page 659.
- Data Service Agents. See [Displaying and Customizing a Data Service Agents Page](#) on page 661.
- Data Service Statistics. See [Displaying Status Statistics and Aggregated Statistics for Data Services](#) on page 669.
- Scheduler Agents. See [Displaying and Customizing a Data Service Agents Page](#) on page 661.
- Web Sessions. See [Web Sessions](#) on page 697.
- Live Console. See [How to Display the Live EDAPRINT Log File](#) on page 760.

Diagnostic

- Savediag - Report a Bug. See [Gathering Diagnostic Information for Customer Support Services](#) on page 794.
- Download. See [Analyzing TIBCO WebFOCUS Reporting Server Activity](#) on page 759 and [Tracing TIBCO WebFOCUS Reporting Server Activity](#) on page 768.

Reporting Server Web Console Appearance

The Reporting Server Web Console Appearance parameters control the look of the Reporting Server browser interface, as well as accessibility and menu bar customization.

To access these settings, go to the *Workspace* page. Click *Settings*, then *Reporting Server Web Console Appearance*.

The Reporting Server Web Console Appearance page opens, as shown in the following image.

The screenshot shows a browser window with four tabs: 'Data Services Agents X', 'New Data Service X', 'Workspace Logs Properties X', and 'Web Console Appearance X'. The active tab is 'Web Console Appearance X'. The page content includes a header 'Menu Bar Appearance' with a small upward arrow icon. Below this, there are two text input fields. The first is labeled 'Caption Of The Custom Link' with a help icon. The second is labeled 'URL Of The Custom Link' with a help icon. At the bottom right, there are two buttons: 'Cancel' and 'Save'.

The following Reporting Server Web Console Appearance settings are saved in the edaserv.cfg file.

Menu Bar Appearance

Caption Of The Custom Link

Specifies the text of a custom link that is shown in the menu bar of the Reporting Server browser interface.

URL Of The Custom Link

Specifies the URL of a custom link that is shown in the menu bar of the Reporting Server browser interface.

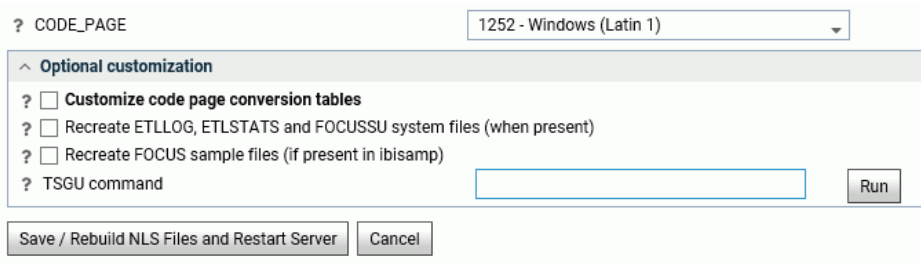
When you have completed your entries, click Save.

Locale Settings

The Reporting Server browser interface provides parameters that define locale-specific date, time, and currency formats as well as parameters for configuring National Language Support.

Reference: NLS Settings

To access these settings, go to the *Workspace* page and click *Settings*. Then point to *Locale* and click *Configuration Wizard*. The NLS Configuration Wizard opens, as shown in the following image.



The following NLS settings are saved in the `nlscfg.err` file:

`CODE_PAGE`

Is the code page of the data sources accessed by the server.

Select a code page from the drop-down list of supported code pages, including Unicode.

The default value is 1252 for Windows and UNIX, and 37 for IBM i (formerly known as i5/OS) and z/OS.

The current value of `CODE_PAGE` is available to applications as the `&FOCCODEPAGE` variable.

Optional customization

`Customize code page conversion tables`

Select this check box to display a list of code pages and descriptions for the possible data sources accessed by the server or the client code page.

Several code pages are always selected by default and cannot be deselected. These are:

- 37 –IBM MF United States
- 137 –U.S. English/Western Europe (Latin 1)
- 437 –U.S. English
- 1047 –IBM MF Open Systems (Latin 1)
- 1252 –Windows Latin; 65001 –Unicode (UTF-8)
- 65001 Unicode (UTF-8)
- 65002 Unicode (UTF-EBCDIC)

Other code pages may also be selected by default, depending on the LANG attribute chosen.

The NLS Configuration Wizard creates conversion tables for these code pages to make them available for subsequent conversion.

Select any additional code pages for which you want tables to be created. For example, if you want to designate a particular code page to be generated in synonyms for flat (fixed) files or VSAM files, you must already have created a code page conversion table using the *Customize code page conversion table* option during NLS configuration.

Recreate ETLLOG, ETLSTATS and FOCUSSU system files

Check this box if you wish to load FOCUS internal files in the server code page (as defined in the CODE_PAGE field).

This is recommended if the compatibility of the new and prior code pages is unknown.

Recreate FOCUS sample files

Check this box if you wish to load the sample files (in the ibisamp application) in the server code page (as defined in the CODE_PAGE field).

TSGU command

TSGU is a utility program to configure NLS settings. It is usually not required, but in the following cases, it must be executed manually.

- To create PDFXTBL.** When TYPE 1 fonts should be included in PDF format file, the TYPE 1 font transcoding table must be generated using the following syntax:

```
PDFX {font-name}
```

- To create PDFYTBL.** When Unicode fonts should be used in PDF format file, the special font information tables (PDFYTBL) must be generated using the following syntax:

```
KTBL {pdfy-name}
```

Complete the configuration, then click *Save/Rebuild NLS Files and Restart Server*. The server creates the necessary transcoding, sorting, and monocasing tables for the selected code page.

Reference: Locale Parameters

The Reporting Server browser interface provides parameters for configuring locale-based date, time, and currency parameters.

To access these settings, go to the *Workspace* page and click *Settings*. Then, click *Locale* and select *Locale (Language, Numbers, Currency, Dates)*. The *Change Core Engine Settings* in the *edasprof.prf* page opens, as shown in the following image.

The following parameter values can be set and stored in the *edasprof.prf* file.

LANGUAGE

Controls the default value of the code page, and the language of server error messages, if available in translation.

Choose a language name or abbreviation from the drop-down list.

The default value is AMENGLISH (American English) on all platforms.

Localized error message files exist for certain languages. If you choose a language for which localized error message files do not exist, the server messages and console appear in English. If you would like the messages and console to appear in a language for which the server has not been localized, you can provide translation yourself. To localize the console, translate the file *wceng.lng* and rename it to *wcXXX.lng* where 'XXX' is the three-letter language code for the server. For error messages, translate all or any relevant **.err* files and append the same three-letter language code to it. All files are located in the `\home\nls` folder.

The current value of LANG is available to applications as the `&FOCLANGCODE` variable. However, it will be blank if it is not configured.

COLLATION

Establishes a binary or case-insensitive collation sequence. The values are:

- CODEPAGE.** Bases the collation sequence on the code page in effect. This is the default value.

- BINARY.** Bases the collation sequence on binary values.
- SRV_CS.** Bases the collation sequence on the LANGUAGE setting, and is case-sensitive.
- SRV_CI.** Bases the collation sequence on the LANGUAGE setting, and is case-insensitive.
- SRV_UNQ.** Bases the collation sequence on the LANGUAGE setting, and is LANGIDUNQ.
- SRV_CSAS.** Bases the collation sequence on the LANGUAGE setting, and is both case-sensitive and accent-sensitive.

When code page files are not modified, CODEPAGE is the same as BINARY, except for Danish, Finnish, German, Norwegian, or Swedish in EBCDIC environment.

The SET COLLATION command has the same effect, and overrides this default setting.

CDN

Defines the punctuation symbols to be used for numbers. Valid values are:

- COMMAS_DOT.** Uses a dot (.) as the decimal separator and comma (,) as the thousands separator. The alias is *OFF*. This is the default value.
- DOTS_COMMA.** Uses a comma as the decimal separator and dot as the thousands separator. The alias is *ON*.
- SPACES_COMMA.** Uses a comma (,) as the decimal separator and space as the thousands separator. The alias is *SPACE*.
- SPACES_DOT.** Uses a dot (.) as the decimal separator and space as the thousands separator. The alias is *SPACEP*.
- QUOTES_COMMA.** Uses a comma (,) as the decimal separator and single quotation mark (') as the thousands separator. The alias is *QUOTE*.
- QUOTES_DOT.** Uses a dot (.) as the decimal separator and single quotation mark (') as the thousands separator. The alias is *QUOTEPE*.

CURRENCY_ISO_CODE

Is a standard three-character currency code such as USD for US Dollars or JPY for Japanese Yen. The default value is *default*, which uses the currency code for the configured language code.

CURRENCY_DISPLAY

Defines the position of the currency symbol relative to a number. The default value is *default*, which uses the position for the format and currency symbol in effect. Valid values are:

- LEFT_FIXED.** The currency symbol is left-justified preceding the number.
- LEFT_FIXED_SPACE.** The currency symbol is left-justified preceding the number, with at least one space between the symbol and the number.
- LEFT_FLOAT.** The currency symbol precedes the number, with no space between them.
- LEFT_FLOAT_SPACE.** The currency symbol precedes the number, with one space between them.
- TRAILING.** The currency symbol follows the number, with no space between them.
- TRAILING_SPACE.** The currency symbol follows the number, with one space between them.

CURRENCY_PRINT_ISO

Defines what will happen when the currency symbol cannot be displayed by the code page in effect. Valid values are:

- DEFAULT.** Replaces the currency symbol with its ISO code, when the symbol cannot be displayed by the code page in effect. This is the default value.
- ALWAYS.** Always replaces the currency symbol with its ISO code.
- NEVER.** Never replaces the currency symbol with its ISO code. If the currency symbol cannot be displayed by the code page in effect, it will not be printed at all.

DATE_ORDER

Defines the order of date components for display. Valid values are:

- DEFAULT.** Respects the original order of date components. This is the default value.
- DMY.** Displays all dates in day/month/year order.
- MDY.** Displays all dates in month/day/year order.
- YMD.** Displays all dates in year/month/day order.

DATE_SEPARATOR

Defines the separator for date components for display. Valid values are:

- DEFAULT.** Respects the separator defined by the USAGE format of the field.
- SLASH.** Uses a slash (/) to separate date components.
- DASH.** Uses a dash (-) to separate date components.
- BLANK.** Uses a blank to separate date components.
- DOT.** Uses a dot (.) to separate date components.
- NONE.** Does not separate date components.

TIME_SEPARATOR

Defines the separator for time components for the &TOD system variable. Valid values are:

- DOT.** Uses a dot (.) to separate time components. This is the default value.
- COLON.** Uses a colon (:) to separate time components.

BUSDAYS

Defines which days are business days for date-time functions. The default value is `_MTWTF_`.

HDAY

Specifies a holiday file for listing dates that are non-business days for use with date functions. The default value is OFF (no holiday file).

WEEKFIRST

Defines the first day of the week for the WEEK and WEEKDAY components of date and date-time functions.

FY_FIRST_MONTH

Defines the first month of the fiscal year. The default value is January.

Reference: Font Support for Arabic and Thai PDF and PostScript Formats

Users that want to output Arabic characters in PDF or PostScript formats can select the Times New Roman font.

Users that want to output Thai characters in PDF or PostScript formats can select the Tahoma font.

Reference: Language and Default Code Pages for Windows and UNIX

This chart provides information generated by the ? LANG SET command from the Windows and UNIX platforms.

Language Code	Language Name	Language Abbreviation	Language ID	Code Page
001	AMENGLISH	AME	en	1252
001	ENGLISH	ENG	en	1252
020	ARABIC	ARB	ar	1256
10351	B-PORTUGUESE	BRA	br	1252
385	CROATIAN	HRV	hr	1250
420	CZECH	CZE	cs	1250
045	DANISH	DAN	da	1252
031	DUTCH	DUT	nl	1252
372	ESTONIAN	EST	et	923
358	FINNISH	FIN	fi	1252
033	FRENCH	FRE	fr	1252
033	FRENCH	FRE	fc	1252
049	GERMAN	GER	de	1252
049	GERMAN	GER	at	1252
030	GREEK	GRE	el	1253
972	HEBREW	HEB	iw	916
036	HUNGARIAN	HUN	hu	1250
039	ITALIAN	ITA	it	1252
081	JAPANESE	JPN	ja	942

Language Code	Language Name	Language Abbreviation	Language ID	Code Page
081	JAPANESE-EUC	JPE	je	10942
082	KOREAN	KOR	ko	949
371	LATVIAN	LVA	lv	1257
370	LITHUANIAN	LTU	lt	1257
047	NORWEGIAN	NOR	no	1252
048	POLISH	POL	pl	1250
351	PORTUGUESE	POR	pt	1252
040	ROMANIAN	ROM	ro	1250
007	RUSSIAN	RUS	ru	1251
085	S-CHINESE	PRC	zh	946
421	SLOVAK	SVK	sk	1250
034	SPANISH	SPA	es	1252
046	SWEDISH	SWE	sv	1252
086	T-CHINESE	ROC	tw	10948
066	THAI	THA	th	874
090	TURKISH	TUR	tr	920
044	UKENGLISH	UKE	uk	1252
800	UNICODE	UCS	uc	65001

Reference: Language and Default Code Pages for IBM i and z/OS

This chart provides information generated by the ? LANG SET command from the IBM i and z/OS platforms.

Language Code	Language Name	Language Abbreviation	Language ID	Code Page
001	AMENGLISH	AME	en	37
001	ENGLISH	ENG	en	37
020	ARABIC	ARB	ar	420
10351	B-PORTUGUESE	BRA	br	37
420	CZECH	CZE	cz	870
045	DANISH	DAN	da	277
031	DUTCH	DUT	nl	37
372	ESTONIAN	EST	et	1112
358	FINNISH	FIN	fi	278
033	FRENCH	FRE	fr	297
033	FRENCH	FRE	fc	297
049	GERMAN	GER	de	273
049	GERMAN-500	GE5	at	500
030	GREEK	GRE	el	875
972	HEBREW	HEB	iw	424
036	HUNGARIAN	HUN	hu	870
039	ITALIAN	ITA	it	280
081	JAPANESE	JPN	ja	939
10081	JAPANESE-930	JPK	jk	930

Language Code	Language Name	Language Abbreviation	Language ID	Code Page
082	KOREAN	KOR	ko	933
371	LATVIAN	LVA	lv	1112
370	LITHUANIAN	LTU	lt	1112
047	NORWEGIAN	NOR	no	277
048	POLISH	POL	pl	870
351	PORTUGUESE	POR	pt	37
040	ROMANIAN	ROM	ro	870
007	RUSSIAN	RUS	ru	1025
085	S-CHINESE	PRC	zh	935
042	SLOVAK	SVK	sk	870
034	SPANISH	SPA	es	284
046	SWEDISH	SWE	sv	278
086	T-CHINESE	ROC	tw	937
066	THAI	THA	th	838
090	TURKISH	TUR	tr	1026
044	UKENGLISH	UKE	uk	285
800	UNICODE	UCS	uc	65002

Displaying Settings for Reporting Server Browser Interface Preferences

To display or change the Reporting Server browser interface preferences, go to the *Workspace* page. Click *Settings*, then *Settings for Reporting Server browser interface Preferences*.

The Settings for Reporting Server browser interface Preferences page opens.

The Settings for Web Console Preferences page opens.

You can select a profile or procedure in which to save these settings. The default profile is `edasprof.prf`. You can set preferences for the following categories.

- Application Preferences.** Allows you to select values for the following options:
 - Show File/Directory Caption Instead of Name. The default value is *No*.
 - Show Files from Nested Applications. The default value is *No*.
 - Show Also Applications not in APPPATH. The default value is *No*.
 - Show Descriptions on Application Tree. The default value is *No*.
 - Show edahome, edaconf, edaprfu, scaroot, edatemp, edalog and foccache. The default value is *No*.
 - Show Files on File Panel. The default value is *Yes*. If you change this to *No*, files will be shown on the application tree when you open an application folder.
 - Restrict Application Ribbon. The default value is *On*.
- Data Assist.** You can select values for the following options:
 - Initial value for *Optimize Load DM Flow Property*. The default value is *On when opening existing or creating new DM Flow*.
 - Auto-run for Synonym Editor (AUTORUN_SE). Enable automatic rerun of Sample data in Synonym Editor on change. Default value is *Run on Change*.
 - Auto-run for New Target in Business View Editor. Default value is *Do Not Run on Change*.
 - Auto-run for Test Query. Default value is *Run on Change with Limited Data*.
 - Auto-run for Custom Copy. Default value is *Do Not Run on Change*.
 - Auto-run for Upload. Default value is *Run on Change*.
 - Auto show data in Synonym Editor tree. Default value is *No*.
 - Auto show data in Test Query tree. Default value is *No*.
 - Auto show data in Custom Copy tree. Default value is *No*.
 - Auto show data in Upload tree. Default value is *Yes*.
 - Enable Persistent Help. The default value is *Yes*.
 - Display Function Strategy. The default value is *Short Description*. You can select *Syntax* to see the function syntax instead of the description.

- Display Name Strategy. The default value is *Title*. You can select *Alias*, *Description*, or *Name*. If you select *Title*, but there is no Title attribute, the field name is displayed.
- Undo Limit. The default value is 50.
- Aggregate Filter Type. The default value is *WHERE_GROUPED*. You can select *WHERE TOTAL*.
- Auto update related nodes. The default value is *No*.
- Use segment name to qualify field references. The default value is *For duplicate fields only*. You can select *Always*.
- Support extended options. Default value is *No*.
- BV_NAMESPACE value (ON/OFF) in Data Assist. The default value is *OFF*.
- Support Business Folders for Cube. The default value is *Yes*.
- Auto-run for Compound/Layout (AUTORUN_COMPOUND). Enable automatic run of Procedures in Compound/Layout on change. Default value is *Run on Change*.
- Use monotree when selecting files in the Data Assist. Default value is *Yes*.
- Automatically redo prepare. Default value is *No*.
- Include keys for default Business View. Default value is *No*.
- Start Get Data Advanced Mode by Default. Default value is *No*.
- Include Fields in Business View. Shows only Business View fields in the Synonym Editor, or all fields. Default value is *BV*.
- Data Assist Joins.** You can select values for the following options.
 - Insert child segments with snowflakes. Default value is *No*.
 - Default Join Type. The default value is *Multiple*.
 - Display full component names (NTM_DSPL_FULL_NAMES). Default value is *Yes*.
 - Use application name to qualify synonym reference (NTM_USE_2PART_NAME. Default value is *Yes*.
 - Use application name to qualify synonym reference. Default value is *Yes*.
 - Show Undo/Redo button with history. Default value is *No*.

Data Assist (Representative Sampling)

- Enable Sampling. Default value is No.
- Application for representative samples. Default value is *foccache*.
- Show sampling information. Default value is Yes.
- Use ETL-TRG-DBMS for sampling. Default value is *No*.
- Prioritize join when creating sampling. Default value is *No*.
- Seed for representative samples. This field is empty, by default.
- Maximum record count for representative samples. This field is empty, by default.

Data Flow Preferences.

- Use application directory name with flow components. Default value is Yes.
- Show Notes. Default value is *No*.
- Calculate source count. Default value is Yes.
- Automatically select all columns. Default value is Yes.
- Add Join Object if needed. Default value is Yes.
- Automatically add join conditions. Default value is Yes.
- Show adapter/application info in the graph. Default value is *No*.
- Automatically map all columns. Default value is *No*.
- Prior to load value (Relational). Default value is *No changes*.
- Prior to load value (Focus). Default value is *No changes*.
- Prior to load value (Flat File). Default value is *No changes*.
- Enable extended data flow support. Default value is *No*.

Index Preferences

- Current Search Engine. Default value is *Lucene*.

Process Flow Preferences.

- Use application directory name with flow components. Default value is Yes.

- Add flows to run as stored procedures. Default value is *No*.
- Show Notes. Default value is *No*.
- Reporting Filter.** You can select values for the following options:
 - Database reads for a filter limit. You can enter a limit for initial database reads.
 - Number of rows in bins for numeric data. Default value is 32.
 - Omit missing (null) values in Data Profiling charts. The default value is *No*.
 - Bins for date unit. Default is *Auto*. You can select *Year*, *Week*, or *Month*.
- Target Defaults.**
 - Default ETL Target DBMS. You can select a configured adapter and connection for target tables in Upload, Quick Copy and Custom Copy.
 - Enable SQL ETL Targets. Default value is *Yes*.
 - Default ETL Target Application. Enter the name of the target application directory for target tables in Upload, Quick Copy and Custom Copy, or click the ellipsis (. . .) to select an application directory.
 - Enable Bulk Load Targets. Default value is *Yes*.
 - Show Load Options prompts. Default value is *Yes*.
 - Upload - prefix synonym name, table name, data set name with user ID. Default value is *No*.

Click **Save** when you have finished entering your preferences.

TIBCO FOCUS Sets and Info

You can generate reports about many aspects of WebFOCUS processing using the Focus Sets and Info option on the Workspace Settings menu.

To access these settings, go to the *Workspace* page. Click *Settings*, then *Focus Sets and Info*.

Core Engine Settings

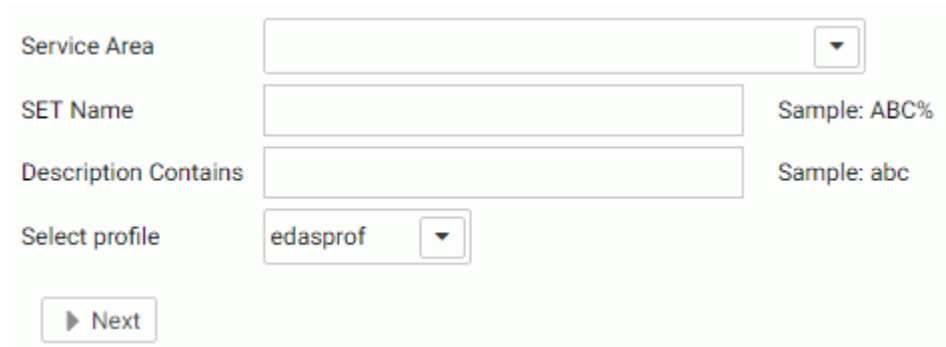
Core engine settings allow you to control a wide range of services, including Dialogue Manager, LOCALE, StyleSheet, Common, TABLE, HLI, Tracing, and Other Sets.

To access the Core Engine settings:

1. Go to the *Workspace* page.

2. Click *Settings*, point to *FOCUS Sets and Info*, then click *Core Engine Settings*.

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



The image shows a web form titled "Core Engine Settings". It contains four input fields: "Service Area" (a dropdown menu), "SET Name" (a text box with a "Sample: ABC%" label), "Description Contains" (a text box with a "Sample: abc" label), and "Select profile" (a dropdown menu with "edasprof" selected). Below these fields is a "Next" button with a right-pointing arrow.

3. Select one or more services from the *Service Area* list.
4. Optionally, enter a string in the *SET Name* field to find specific settings, using % as a wildcard.

For example, to find settings containing DATE, enter DATE%. All settings are uppercase and the *SET Name* field is case-sensitive.
5. Optionally, enter a string in the *Description Contains* field to search for settings by their descriptions.
6. Optionally, select a profile from the *Select profile* drop-down menu.
7. Click *Next*.

The Core Engine Settings pane opens.

8. Make your changes and click *Save*.

Example: Controlling Error Message Processing With the EMGSRV and EMGLIMIT Parameters

To control error message routing, use the EMGSRV parameter. You can route messages to the normal output device (the default), suppress error message output, route messages to a file named emgfile.ftm in your current directory, route messages to the edaprint.log file, or route them to both edaprint and the normal output device. The EMGLIMIT parameter limits the number of error messages written to the edaprint log file.

You can set these parameters using syntax in a procedure or using the core engine settings page in the Reporting Server browser interface.

The syntax for the EMGSRV parameter is

```
SET EMGSRV = {ON|OFF|FILE|EDAPRINT|EDAPRINT/ON}
```

where:

ON

Routes messages to the normal output device (no error message re-routing). This is the default value.

OFF

Suppresses error message output.

FILE

Routes error messages to emgfile.ftm in the current directory. If needed, use an APP COPYFILE command to save the emgfile.ftm file to other locations.

EDAPRINT

Routes error messages to the edaprint log file.

The EDAPRINT option allows Workspace Event Routing to be used to detect messages and trigger processes and email alerts.

EDAPRINT/ON

Routes error messages to both the normal output device and the edaprint log file.

The syntax for the EMGLIMIT parameter is

```
SET EMGLIMIT = n
```

where:

n

Is the maximum number of error messages to be written to the log. The default value, zero (0), specifies an unlimited number of messages.

To use the Reporting Server browser interface core engine settings page for these parameters, enter EMG% in the SET Name text box and click *Next*.

Enter a number in the EMGLIMIT text box to limit the number of messages written to the log.

Select the FILE option for EMGSRV to route error messages to the file emgfile.ftm in your current directory, as shown in the following image.

The screenshot shows a configuration window titled "Error / Warning Control". It contains two settings:

- EMGSRV:** A dropdown menu currently showing "FILE - Route to emgfile.ftm (current directory)". To its right is the text "General message control (ON=Normal, OFF=Suppress All)".
- EMGLIMIT:** A text input box that is currently empty. To its right is the text "Set limit of Error Messages Printed (Default=0 (Unlimited))".

At the bottom of the window, there are three buttons: "Back", "Save", and "Cancel".

Click **Save** to save these settings.

Core Engine System Variables

Core Engine System variables control the functionality of the server.

The variables fall into the following categories:

- Display Information
- Operating System
- Release Information
- National Language Support
- File Information
- Statistical Variables
- Date/Time Variables
- Security Information
- Miscellaneous System Variables

To view Core System variables:

1. Go to the *Workspace* page.
2. Click *Settings*, point to *FOCUS Sets and Info*, and click *Core Engine System Variables*.

The Core System Variable pane opens, as shown in the following image.



The image shows a configuration interface for Core Engine System Variables. It consists of three main input fields stacked vertically. The first field is labeled 'Service Area' and is a dropdown menu. The second field is labeled 'Variable Name' and is a text input field; to its right is the text 'Sample: ABC%'. The third field is labeled 'Description Contains' and is also a text input field; to its right is the text 'Sample: abc'. Below these fields is a button labeled 'Next' with a right-pointing arrow.

3. Select one or more services from the *Service Area* list.
4. Optionally, enter a string in the *Variable Name* field to find specific variables, using % as a wildcard. Variables are all uppercase and all begin with an ampersand (&).
5. Optionally, enter a string in the *Description Contains* field to search for variables by their descriptions.

- Click *Next*.

The Core System Variables are displayed.

Core Engine Global Variables

You can add global variables to a profile on the server to deliver default values or execute SET commands.

To create Core Engine global variables:

- Go to the *Workspace* page.
- Click *Settings*, point to *FOCUS Sets and Info*, and click *Core Engine Global Variables*.

The Core Engine Global Variable pane opens, as shown in the following image.

The image shows a user interface for adding global variables. It features a 'Select profile' dropdown menu with 'edasprof' selected. Below the menu is a large, light gray 'Add New' button.

- Optionally, select a profile from the drop-down menu. The default is edasprof.
- Click *Add New*.
- On the Core Engine Global Variables pane, select *SET* or *DEFAULT* from the *Command* drop-down menu.
- Enter a name and expression for the variable in the corresponding fields, as shown in the following image.

The image shows the 'Add New' form for a global variable. At the top, there is a 'Select profile' dropdown menu with 'edasprof' selected. Below it is a note: 'Note: The Value is a constant value for DEFAULT and an expression for SET command'. The form has three main fields: 'Command' (a dropdown menu set to 'DEFAULT'), 'Name' (a text input field containing 'LIMIT'), and 'Expression' (a text input field containing '100'). At the bottom of the form are two buttons: 'Add New' and 'Save'.

To add an additional global variable, click *Add New*.

7. Click Save.

Displaying HOLD Keywords

You can display the keywords available for generating different types of HOLD files.

1. Go to the *Workspace* page.
2. Click *Settings*, point to *FOCUS Sets and Info*, and click *HOLD Keywords*.

The *HOLD Keywords* page opens, as shown in the following image.



The screenshot shows a configuration form with the following elements:

- Service Area:** A dropdown menu.
- Keyword:** A text input field with a "Sample: ABC%" label to its right.
- Description Contains:** A text input field with a "Sample: abc" label to its right.
- Next:** A button with a right-pointing arrow and the text "Next".

3. Select the following parameters.

- Optionally, select one or more Service Areas from the *Service Area* drop-down list.

The areas are:

- Keywords specific to Relational DBMS formats.
 - Keywords for Extended Bulk Load.
 - Keywords specific to format DFIX.
 - Keywords specific to format XML.
 - Keywords specific to Excel formats.
 - Keywords specific to stylesheet formats.
 - Other format-specific keywords.
 - Keywords to specify output for the data.
 - Miscellaneous keywords.
 - Obsolete keywords.
- Optionally, enter a keyword or keyword pattern in the *Keyword* text box.

- Optionally, enter a description or description pattern in the *Description Contains* text box for the keywords you want to see.
- Click *Next*.

The HOLD Keywords Report is displayed, as shown in the following image.

^ Keywords specific to FORMAT DFIX	
DELIMITER	Specifies delimiter for HOLD FORMAT DFIX
ENCLOSURE	Specifies enclosure character(s) for alphanumeric data values for HOLD FORMAT DFIX
HEADER	Specifies if header has to be included into data file for HOLD FORMAT DFIX
NO	Allowed value for HEADER and PRESERVESPACE for FORMAT DFIX
PRESERVESPACE	Does not trim trailing spaces in data for FORMAT DFIX
RDELIMITER	Specifies end-of-row character(s) for HOLD FORMAT DFIX
YES	Allowed value for HEADER and PRESERVESPACE for FORMAT DFIX
^ Keywords specific to Excel formats	
CACHEFIELDS	Specifies cache fields for FORMAT EXL2K PIVOT
FORMULA	FORMAT EXL2K FORMULA specifies that the report is for Excel 2000 with formulas
PAGEFIELDS	Specifies page fields for FORMAT EXL2K PIVOT
PIVOT	FORMAT EXL2K PIVOT specifies that the report is for Excel 2000 with a pivot table.
SHEETNUM	See SHEETNUMBER
SHEETNUMBER	Specifies the number of the worksheet being replaced in the template for FORMAT EXL2K.
TEMPLATE	Specifies the name of the Excel template file (workbook)

[◀ Back](#)

Displaying FOCUS Limits

You can display limits for WebFOCUS requests and components, as well as limits that have been removed.

1. Go to the *Workspace* page.
2. Click *Settings*, point to *FOCUS Sets and Info*, and click *Focus Limits*.

The *Focus Limits* page opens, as shown in the following image.



The screenshot shows a configuration form with the following elements:

- Service Area:** A dropdown menu.
- Limit:** A text input field with a label "Sample: ABC%" to its right.
- Description Contains:** A text input field with a label "Sample: abc" to its right.
- Next:** A button with a right-pointing arrow and the text "Next" below it.

3. Select the following parameters.

- Optionally, select one or more service areas from the *Service Area* drop-down list.

You can select from the following service areas:

- Limits specific to Master Files.
 - Limits specific to TABLE.
 - Limits specific to JOIN and COMBINE.
 - Limits specific to I/O.
 - Limits specific to MODIFY.
 - Limits specific to CREATE SYNONYM.
 - Miscellaneous limits.
 - Removed limits.
- Optionally, enter a limit name or pattern in the *Limit* text box.
 - Optionally, enter a description or description pattern in the *Description Contains* text box.
 - Click *Next*.

The FOCUS Limits Description report displays, as shown in the following image.

^ Limits specific to Master File Description		
MAX_AIND_LEN	256	Maximum length of index for alpha field in focus MFD
MAX_FOCINDX_LEN	12	Size of a FOCUS index field name
MAX_FOC_SEG_LEN	3968	Maximum length of a FOCUS segment
MAX_XFOC_SEG_LEN	16284	Maximum length of a XFOC segment
NTM_LEVCM_MAXSEG	1024	Possible number of segments
NTM_LEVCM_MAXSEG_XFOC	512	Possible number of segments in XFOCUS
NTM_MAX_ACTUAL_LEN	14	Length of ACTUAL
NTM_MAX_DATEPATTERN_LEN	64	Length of a date pattern
NTM_MAX_FLD_NAME_SIZE	512	Size of a field name/alias
NTM_MAX_FOC_SEG_NAME_SIZE	8	Size of a FOCUS segment name
NTM_MAX_FORMAT_LEN	8	Length of FORMAT
NTM_MAX_HELPMSG_LEN	78	Length of a help message
NTM_MAX_LABEL_LEN	128	Label length
NTM_MAX_QUAL_NAME_SIZE	1090	Size of a qualified field name
NTM_MAX_SEG_NAME_SIZE	64	Size of a segment name
NTM_MAX_SUFFIX_LEN	8	Suffix length
NTM_MAX_TABLE_NAME_SIZE	128	Size of a SQL table name including owner
NTM_MAX_TITLE_LEN	512	Title length
NTM_MODIFY_MAX_FLD	32767	Number of fields in MODIFY/SCAN
NTM_NTOLAP_NAME_SIZE	143	Max length of Dimension/Hierarchy/Measure Group name
NTM_TEXT_FIELD_LIM	64	Number of TEXT fields in one MFD
NTO_CAPTION_LENGTH	180	Max length of Dimension/Hierarchy/Measure Group/KPI caption

◀ Back

Displaying Error Messages

You can display all or different categories of error messages.

1. Go to the *Workspace* page.
2. Click *Settings*, point to *FOCUS Sets and Info*, and click *Error Message Report*.

The *Filter Error Message Report* page opens, as shown in the following image.

Error File

Show Warnings Only

Show Unused/Reserved Only

Text contains

Show Report

3. Select the following parameters.

- Optionally, select one or more error files from the *Error File* drop-down list. Each error file on the list shows the range of error numbers contained in the file.
- Optionally, check the *Show Warnings Only* check box to limit the report to warning messages.
- Optionally, check the *Show Unused/Reserved Only* check box to control whether only unused or reserved messages display in the report.
- Optionally, enter text in the *Text Contains* text box for the messages you want to see.
- Click *Show Report*.

The Error Message report displays in a new tab, as shown in the following image.

Error Number	Error Text
201	(FOC201) INTERRUPT. DIVISION BY ZERO
202	(FOC202) INTERRUPT. FLOATING VALUE OVERFLOW
203	(FOC203) INTERRUPT. FLOATING VALUE UNDERFLOW
204	(FOC204) TOO MANY INTERRUPTS. RUN TERMINATED
205	(FOC205) THE DESCRIPTION CANNOT BE FOUND FOR FILE NAMED: %1
206	(FOC206) INSUFFICIENT MEMORY IS AVAILABLE FOR PROGRAM/CACHE
207	(FOC207) ERROR IN THE FORMAT DEFINITION OF FIELD: %1
208	(FOC208) REENTRANCY ERROR. MODULE DOES NOT CONTAIN CSECT: %1
209	(FOC209) THE DATA VALUE EXCEEDS ITS LENGTH SPECIFICATION: %1
210	(FOC210) THE DATA VALUE HAS A FORMAT ERROR: %1
211	(FOC211) THE PARENT SEGMENT NAME IS NOT FOUND IN THE MASTER: %1
212	(FOC212) THE FILE DESCRIPTION CONTAINS AN ERROR: %1
213	(FOC213) THE RECORD CONTAINS A FIELDNAME WHICH IS NOT RECOGNIZED: %1
214	(FOC214) NUMBER OF SEGMENTS OR SEGMENTS+INDEXES EXCEEDS THE MAXIMUM: %1
215	(FOC215) THE FIRST SEGMENT IS SPECIFIED WITH A NON-BLANK PARENT: %1
216	(FOC216) THE TERMINATOR OF '\$' MISSING FROM THE END OF THE RECORD
217	(FOC217) THE SEGMENT NAME IS A DUPLICATE AND NOT THE SAME AS PREVIOUS: %1

Displaying the Currency Report

The currency report lists currency names, ISO codes, currency symbols, currency symbols hexadecimal values, currency symbol character lengths, and currency symbol byte lengths.

To display the Currency Report, go to the *Workspace* page, click *Settings*, point to *FOCUS Sets and Info*, and click *Currency Report*. The Currency Report page opens, as shown in the following image.

Currency name	ISO code	Symbol	Symbol hexadecimal	Char-length	Byte-length
USA dollar	USD	\$	24	1	1
Eurozone euro	EUR	€	80	1	1
Japanese yen	JPY	¥	a5	1	1
United Kingdom pound sterling	GBP	£	a3	1	1
United Arab Emirates dirham	AED			0	0
Afghan afghani	AFN	Afs	41, 66, 73	3	3
Albanian lek	ALL	Lek	4c, 65, 6b	3	3
Netherlands Antillean guilder	ANG	f	83	1	1
Angolan kwanza	AOA	Kz	4b, 7a	2	2
Armenian dram	ARM			0	0
Argentine peso	ARS	\$	24	1	1
Australian dollar	AUD	\$	24	1	1
Aruban florin	AWG	Afl.	41, 66, 6c, 2e	4	4
Azerbaijani manat	AZN			0	0
Bosnia and Herzegovina convertible mark	BAM	KM	4b, 4d	2	2
Barbados dollar	BBD	\$	24	1	1
Bangladeshi taka	BDT	Tk	54, 6b	2	2

Determining Function Parameters

You can determine the parameters for a particular function from the Function Parameters Report. The report also includes a brief description of the function.

To display the Function Parameters Report:

1. Go to the *Workspace* page.
2. Click *Settings*, point to *FOCUS Sets and Info*, and click *Function Parameters Report*.

The *IBI Functions Report* page opens, as shown in the following image.

The screenshot shows a configuration panel for the IBI Functions Report. At the top, there is a question mark icon followed by the text "Select function category". To the right of this text is a dropdown menu currently set to "Recommended". Below this is another dropdown menu showing "Aggregation Operation, Analyt". Underneath, there is a section with a question mark icon and a checkbox labeled "Filter by function name AND/OR description". Below this are three more checkboxes: "Show Function Parameters" (which is checked), "Show Function Example", and "Show Report in PDF". At the bottom of the panel is a button labeled "Show Report".

3. Select the following parameters.

- Optionally, select one or more function categories from the *Function Category* drop-down list. The choices are Recommended, Legacy, and All.
- All of the functions in the category you chose are selected by default. Optionally, deselect one or more function types from the drop-down list, if you do not want to see those function types.
- Optionally, click *Filter by function name AND/OR description* to provide a name or pattern for the function names and/or descriptions you want to see. as shown in the following image.

This screenshot shows the configuration for the "Filter by function name AND/OR description" option. It includes a warning: "Warning: search is case insensitive". There are two input fields: "Function name like:" and "Function description contains:". The "Function name like:" field has a sample text: "Sample: abc%. Wildcard charactes '%' and '_' is escape character". The "Function description contains:" field has a sample text: "Sample: abc". There are radio buttons for "AND" (which is selected) and "OR".

- Optionally, click *Show Function Example* to see an example of using the function.
- Optionally, click *Show Report in PDF* to generate a PDF version of the report.

4. Click *Show Report*.

The Function Parameters Report is displayed, as shown in the following image.

Function Category: Format Conversion		
Name	Type	Description
COMPACTFORMAT(input_field)	Alpha	Represent a numeric field in a suitable compacted form.
input_field	Numeric	Is the name of a numeric field.
DIGIT\$(number, out_length)	Variable Alpha	Create a string from numeric. Evaluate numeric expression as a whole number with fraction truncated and returns last N digits of its absolute value.
number	Integer	Is any numeric value whose whole absolute value will be put in output string.
out_length	Integer Constant	Is the number of output rightmost digits. It is a whole number from 1 to 10.
DT_FORMAT(date, date_format)	Variable Alpha	Date or Datetime as a character string. Return a character string after the date conversion according to the date_format provided.
date	Numeric, Date-Time, Date	Is date or datetime type field. Date may be given in any of the standard date format types; datetime in any datetime format type.
date_format	Alpha Constant	Is a literal in single quotation marks which represents the format for conversion that fits the date format type.
FPRINT(value, 'out_format')	Alpha	Print using format. Return a character string after conversion of input data according to the format provided.
value	any type	Is the value to be converted.
'out_format'	Fixed Alpha	Is the output format.
HEXTYPE(in_value)	Variable Alpha	Return hex value. Return the hexadecimal view of input value.
in_value	any type	Any data format field, constant, expression.
TO_CHAR(value)	Alpha	Convert numeric to character string. Convert numeric data to a character string.
value	Numeric, Date-Time, Date	Is any numeric field or expression.
TO_INTEGER(string)	Integer	Convert to integer. Convert a string to an integer value.
string	Alpha	Is any type of character string to be converted into its integer value.
TO_NUMBER(string)	Numeric	Convert to number. Convert a string into the numeric format most appropriate to the context.
string	Alpha	Is any type of character string to be converted into its double precision floating point value.

Workspace SET

The Workspace SET option allows you to configure SMTP Server settings, Profile settings, Application settings, and Miscellaneous settings.

Reference: SMTP Server Settings

You can specify an IP address for the email server and a sender address. If the server is secured, you can specify the credentials for connecting to the server. You can also test your email connection from the E-mail Settings pane.

To access these settings, go to the *Workspace* page. Click *Settings*, point to *Workspace SET*, and click *SMTP Server Settings*.

The E-Mail SMTP Server page opens, as shown in the following image.

Note: If an antivirus program is running on this computer, please make sure that the processes tscm3.exe, edapgwv.exe and edaplog.exe are excluded from the processes which are prevented from sending email.

? smtp_host	<input type="text"/>
? port	<input type="text" value="25"/>
? smtp_user	<input type="text"/>
? smtp_password	<input type="password"/>
? sender_email	<input type="text"/>
? server_admin_email	To <input type="text"/> <input type="button" value="v"/> Use semicolon to separate emails, user IDs, procedures.
<input type="button" value="Save and Restart Server"/> <input type="button" value="Send Test E-Mail"/>	

The following settings are saved in the `odin.cfg` file:

`smtp_host`

Indicates the IP address of the SMTP email server using the notation that is supported for your IP version (for example, for IPv4, use base 256 standard dot notation. When IPv6 is supported on a platform, and enabled, you can use either 256 standard dot notation or IPv6 notation.)

`port`

Indicates the TCP port number that a client is connecting to. The default value is 25.

`smtp_user`

Is the user ID to authorize access to the SMTP server.

`smtp_password`

Is the password for the `smtp_user`.

`sender_email`

Is the sender email address (the address that appears in the *From* field in emails).

`server_admin_email`

Is a list that can consist of email addresses, user IDs and procedure names (for example, `%myfex.fex`), each on a separate line, with no delimiters.

Note: If email send operations fail, verify that the SMTP host is active and/or that a virus scanner is not blocking `edachkup`.

When you have completed your entries, optionally click the *Send Test E-mail* button to test your connection, then click *Save and Restart Server*.

Reference: Profile Settings

You can control aspects of profile execution from the Profile Settings pane.

To access these settings, go to the *Workspace* page. Click *Settings*, point to *Workspace SET*, and click *Profile Settings*.

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

The following Profile Settings are saved in the edaserv.cfg file:

`profile_parameter`

By passing parameters to profiles, you can submit different session parameters at different connection times. When this parameter is activated, any unresolved variables are prompted for during your initial logon to the Reporting Server browser interface.

Available settings are *y* and *n*.

If the server, user, or group profiles contain any unresolved variables, set the profile parameter keyword to *y*. The default setting is *n*.

`connections_autoprompt`

Enables prompting for user credentials on running a procedure and upload. When this is set to *y*, the server will:

- Add an additional parameter to the connection string for SQL Adapters. You can configure this on the Reporting Server browser interface Adapter Connection Properties page by setting *Auto Prompt Connection Credentials* to *y*. The CONNECTION_ATTRIBUTES_EXT statement in the server profile will be updated with "auto_prompt": "y".
- Add the "Run As" option for procedures on the Reporting Server browser interface. This will return all connections from all Master Files used in procedures that are configured with auto_prompt": "y".

`profile_no_domain *`

In the Windows Server, to control whether the domain name is included in or excluded from a user profile name, you can choose one of the following.

n to add the domain name to the profile name (for example, user ibi/edaport becomes ibi_edaport.prf). *N* is the default value.

y to omit the domain name from the profile name (for example, user ibi/edaport becomes edaport.prf). This setting is required for compatibility with earlier releases of the Server.

Note: The *y* setting can be safely employed if a single domain is in use and there are no local user IDs that could interfere with domain IDs. If there is potential for interference between local user IDs and domain IDs that are connected to the server, the *n* setting is required to ensure that each user has a unique profile name.

Note that server must be restarted when this parameter is changed.

When you have completed your entries, click Save.

Reference: Application Settings

When you want to be able to reuse data within the same browser session, you can store the data in the form of a HOLD, SAVE, or SAVB file in the foccache directory, which is automatically created when the connection to the server is established. This becomes the first directory in the application path.

As long as the browser session remains active, the stored files remain in the foccache directory and can be referenced in requests using standard two-part names. For example, the first request below creates a HOLD file in foccache, which is referenced by the second request:

```
TABLE FILE CAR
PRINT CAR
ON TABLE HOLD AS FOCCACHE/myfile
END
```

```
TABLE FILE FOCCACHE/myfile
PRINT *
END
```

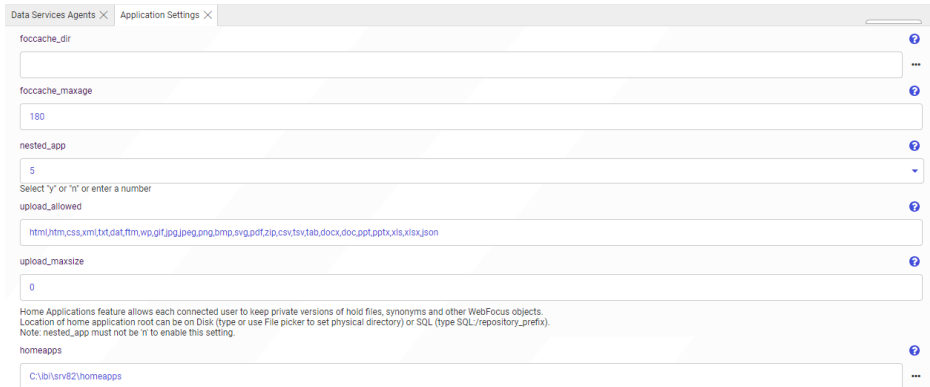
When the maximum foccache retention time is reached (180 minutes by default), the foccache directory is removed by the server and can no longer be referenced.

This feature provides an important mechanism for caching and running multiple instances of WebFOCUS active and drill-down reports. You can use any WebFOCUS graphical tool to create the request and generate a report.

Note: Each time a WebFOCUS user logs into the Managed Reporting (MR) environment, a new session with its own foccache directory is created. Files that were stored in a prior foccache directory are no longer accessible. When the user logs out of MR the foccache directory used within that MR session is no longer accessible.

To access these settings, go to the *Workspace* page. Click *Settings*, point to *Workspace SET*, and click *Application Settings*.

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



The following Application Settings are saved in the `edaserve.cfg` file.

`focache_dir`

Defines the location of the foccache directory. The default is `edatemp/focache/...`

`focache_maxage`

Defines the maximum number of minutes that a foccache directory will last. The default is 180 minutes. When the foccache directory expires, the session also expires.

0 means unlimited (that is, the directory will remain active as long as the session lasts or the server is running).

`nested_app`

For applications containing subfolders or multiple-level subfolder trees, this parameter controls whether subfolders will be automatically added to the application path.

The valid values are:

- A positive integer - Enable nested application, and only a specific number of subfolder levels in the active application path.
- n* - Disables nested applications. Only applications explicitly in the APP PATH will be used.
- y* - Enables nested applications. The entire application subfolder tree under each application in the APP PATH is included in the active application path.

The default value is 5.

Note: For z/OS servers, this setting is only applicable to directory-style applications, it is not applicable to PDS-style applications nor to applications mapped as a collection of ddnames.

`upload_allowed`

Lists the extensions allowed for upload. You can remove extensions from the list to prevent files with those extensions from being uploaded.

`upload_maxsize`

Enables you to change the maximum file size in kilobytes supported for upload. The default value, zero (0), allows files of any size to be uploaded.

If a maximum size is set, attempting to upload a larger file will generate the following message, and the upload will terminate.

`The size of the uploaded file exceeds the limit of nnnn kilobytes`

`homeapps`

If `nested_app` is set to `y` or a number, `homeapps` defines the location of the home application root directory. With `homeapps` defined, each connected user will see a *My Home* applications folder on the Application tree. Users can create new applications (home applications) in this folder. These applications will be visible and available only to the connected user and the server administrator. The user will have full access to his home applications.

When you have completed your entries, click *Save and Restart Server*.

Reference: Miscellaneous Settings

Miscellaneous Settings allow you to set additional Workspace operating parameters, including automatic recovery, transaction coordination, soft and hard kill delays, maximum session log rows, time window for running response time averages, and proxy server settings.

To access these settings, go to the *Workspace* page. Click *Settings*, point to *Workspace SET*, and click *Miscellaneous Settings*.

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

Changing values with an asterisk (*) requires restarting the server

? automatic_recovery	y
? transaction_coordination_mode *	n
? tscom3_kill_delay1	16
? tscom3_kill_delay2	32
? crashed_cleanup *	n
? stats_runningavg *	600
? htmlencode	n
? proxy_server *	
? proxy_port *	
? proxy_user *	
? proxy_password *	
? outbound_ssl_certificate_file *	
? outbound_ssl_certificate_passphrase *	
? outbound_ssl_certificate_label *	
? outbound_ssl_certificate_keyfile *	
? outbound_ssl_auth_server *	n
? seslog_max_lines	1000
? popup_timer	5

Save Cancel

The following Miscellaneous Settings are saved in the edaserv.cfg file:

`automatic_recovery`

Determines whether the server will automatically recover crashed listeners and/or special services. The default is *n*.

`tscom3_kill_delay1`

Defines the time (in seconds) allowed for an agent to stop before soft kill is sent. The default is 16.

`tscom3_kill_delay2`

Defines the time (in seconds) allowed for a soft kill to work before a hard kill is sent. The default is 32.

`crashed_cleanup`

By default, crashed agents are kept in the agents table for diagnostics purposes, and are only removed when they are explicitly killed. When this is set to *y*, crashed agents are automatically removed from the agents list when the Workspace Manager detects a crash. The default value is *n*.

`crashed_max_dump`

If `crashed_cleanup` is set to `y`, this determines the number of memory dump files retained for later diagnosis.

`stats_runningavg`

Defines the time window in seconds for the calculation of running response time averages. The default is 600.

`htmlencode`

Specifies whether to encode HTML tags in data as plain text in order to prevent attacks by inserting executable code into data. Valid values are `y` and `n`. The default value is `n`.

`proxy_server`

Defines a proxy server, which acts as an intermediary on requests for resources from other servers.

`proxy_port`

Defines the port number of the proxy server.

`proxy_user`

Defines a user ID for the proxy server.

`proxy_password`

Defines a password for the user ID connecting to the proxy server.

`outbound_ssl_certificate_file`

Required for SSL when using self-signed certificates on i5 systems.

`outbound_ssl_certificate_passphrase`

Required for SSL when using self-signed certificates on i5 systems.

`outbound_ssl_certificate_label`

Required for SSL when using self-signed certificates on i5 systems.

`outbound_ssl_certificate_keyfile`

Required for SSL when using self-signed certificates on i5 systems.

`outbound_ssl_auth_server`

Required for SSL when using self-signed certificates on i5 systems.

`seslog_max_lines`

Defines the upper limit for the number of rows in a session log. When the limit is exceeded, only the latest records will be kept in the session log and the number of records will equal the limit. The default is 1000.

`popup_timer`

Defines the time in seconds to wait before the popup timer appears. The default is 5.

z/OS Specific Parameters`tmpalloc_primary``tmpalloc_secondary`

Allocation (in CYLs) for temporary Reporting Server browser interface files allocated by the HTTP listener. The default is 5 and 50, respectively.

`submit_opt`

Customizes the call of the IKJEFF10 exit when submitting jobs with the DYNAM SUBMIT command.

`smf_recno`

Defines the SMF record number that the server will use to record utilization statistics.

`smf_subtype`

Defines which pairs of SMF subtype records are cut. The options are:

- all.** Cut all records. This is the default.
- logon.** Cut logon/logoff subtype only.
- query.** Cut query begin/query end subtype only.

When you have completed your entries, click Save.

Procedure: How to Customize the IKJEFF10 Exit for DYNAM SUBMIT From the Reporting Server browser interface (z/OS)

To change how the server calls the IKJEFF10 exit, set the `submit_opt` parameter as follows:

1. Go to the *Workspace* page.
2. Click *Settings*, point to *Workspace SET*, then click *Miscellaneous Settings*.
The Miscellaneous Settings page opens.
3. For the `submit_opt` option, assign a separate value (0 or 1) in each of the three drop-down lists. The default is 000.

The first drop-down list enables the call to IKJEFF10 from DYNAM SUBMIT. By default, the call is disabled.

0

Does not call IKJEFF10. This is the default setting.

1

Calls IKJEFF10.

The second drop-down list specifies that SUBMIT is to call IKJEFF10 in problem state, not in supervisor state. The IBM specification for IKJEFF10 requires supervisor state. If the version of IKJEFF10 in use at your site is capable of running in problem state, this option may be changed, thus allowing DYNAM SUBMIT to be used with IKJEFF10 when the server is not running APF-authorized. When you set:

0

Supervisor state is required for calling IKJEFF10. This is the default setting.

1

Supervisor state is not required to call IKJEFF10.

The third drop-down list allows SUBMIT to continue in the event that an error occurs while processing an in-stream data set. In-stream data sets are created by the following statements in the job:

```
//ddname DD *  
//ddname DD DATA  
/*XMIT  
//name XMIT  
/*DATASET
```

If an error is found in one of the previous statements (for example, a syntax error in the DLM parameter), the SUBMIT operation is normally cancelled. This option allows submission of the job to continue, although in such a case the results are unpredictable since the JCL contained within the in-stream data set may be erroneously interpreted as part of the job being submitted.

0

Terminates SUBMIT on in-stream data set error. This is the default setting.

1

Does not terminate SUBMIT on in-stream data set error.

4. Click Save.

Your entries are registered in the edaserve.cfg file.

Configuring Geographic Information

Geographic roles have been unified to access 58 world administrative boundaries down to the postal code level, out of the box.

The list of geographic roles has been simplified and unified by adding hierarchical levels of keys in the server geographic configuration file. Your request that provides the geography dimension must provide the hierarchical keys needed to identify each target geography. For example, if your geographic role is City, you must also provide values for State and Country.

The following is the list of unified geographic roles is available for generating maps:

- Point of Interest.
- City.
- Continent.
- Country.
- County.
- Geometry line.
- Postal code.
- State.

The server has a geographic configuration editor for customizing the list of geographic roles, basemaps, reference layers, and demographic layers. A customized geographic role can reference either an uploaded shapefile or an existing Esri Feature Layer.

Note:

- You must configure the Adapter for ESRI ArcGIS in order to use many of the geographic features. For more information on configuring the adapter, see the *TIBCO WebFOCUS® Adapter Administration* manual.
- Many names of geographic locations include special characters. In order to map them, the server must be configured for Unicode.

Reference: Using Geographic Administrative Boundaries

Administrative boundaries for the United States are automatically included by Esri.

The following table lists the Level 1 countries and their Level 1 and Level 2 administrative levels. For some countries, no additional administrative levels are added with a second level license.

Country	Postal Code	Level 1 Administrative Level (Number of Values)	Level 2 Administrative Level (Number of Values)
Albania	2 - digit	Qarke (12)	Bashkia (61)
Andorra	5 - digit	Parroquies (7)	
Argentina	5 - digit	Departamentos (525)	
Australia	4 - digit	Local Government Areas (569)	
Austria	4 - digit	Gemeinden (2120)	
Belarus	6 - digit	Woblaszi (7)	Raioni (129)
Belgium	4 - digit	Communes (589)	
Bosnia and Herzegovina	5 - digit	Kantone (18)	Opstine (142)
Brazil	5 - digit	Municipios (5565)	
Canada	3 - digit	Census Divisions (293)	Census Subdivisions (5243)
Chile	2 - digit	Provincias (54)	
China	2 - digit	Counties City Districts (2855)	
Costa Rica	2 - digit	Cantones (81)	Distritos (473)
Croatia	2 - digit	Zupanije (21)	Opcine (556)
Cyprus	2 - digit	Eparchia (6)	Dimoi (615)

Country	Postal Code	Level 1 Administrative Level (Number of Values)	Level 2 Administrative Level (Number of Values)
Czech Republic	2 - digit	Orp (206)	Obce (6253)
Denmark	4 - digit	Kommuner (99)	
Estonia	5 - digit	Maakonnad (15)	Vallad (213)
Finland	5 - digit	Seutukunnat (70)	Kunnat (317)
France	5 - digit	Cantons (1972)	Communes (36571)
Germany	5 - digit	Municipalities	
Gibraltar	3/4/5 - digit		
Greece	5 - digit	Nomoi (75)	Dimoi (326)
Holy See (Vatican City State)	5 - digit		
Hong Kong	6 - digit	Shih Zizhiqu (18)	
Hungary	2 - digit	Kistersegek (175)	Telepulesek (3177)
Iceland	2 - digit	Landsvaedi (8)	Sveitarfelog (74)
India	6 - digit	Districts (640)	Subdistricts (5783)
Italy	5 - digit	Provincias (110)	Comuni (8047)
Japan	3 - digit	Prefectures (47)	Municipalities (1901)
Kenya	1 - digit	Wilaya (47)	Districts (158)
Korea, Republic of	3 - digit	Si Gun (230)	
Latvia	2 - digit	Rajons (33)	Novadi (119)
Liechtenstein	4 - digit	Landschaften (2)	Gemeinden (11)

Country	Postal Code	Level 1 Administrative Level (Number of Values)	Level 2 Administrative Level (Number of Values)
Lithuania	2 - digit	Apskritis (10)	Savivaldybes (60)
Luxembourg	4 - digit	Cantons (12)	Communes (105)
Macedonia, the former Yugoslav Republic of	2 - digit	Statisticki Regioni (8)	Opstini (80)
Malta	3 - digit	Districts (6)	Local Councils (68)
Mexico	3 - digit	Estados (32)	
Monaco	5 - digit	Quartiers (9)	
Montenegro	2 - digit	Opstina (21)	
Netherlands	4 - digit	Gemeenten (380)	
New Zealand	4 - digit	Territorial Authorities (68)	Area Units (1911)
Norway	4 - digit	Okonomisk regioner (89)	Kommuner (428)
Poland	5 - digit	Gminy (2479)	Gminy Miasta (3119)
Portugal	4 - digit	Concelhos (308)	Freguesias (3092)
San Marino	5 - digit	Castelli (9)	
Serbia	5 - digit	National (1)	
Singapore	2 - digit	Regions (5)	Planning Areas (55)
Slovakia	2 - digit	Okresy (79)	Obce (2927)
Slovenia	2 - digit	Statisticne Regije (12)	Obcina (212)

Country	Postal Code	Level 1 Administrative Level (Number of Values)	Level 2 Administrative Level (Number of Values)
South Africa	2-digit	Municipalities (234)	
Spain	5 - digit	Provincias (52)	Municipios (8122)
Sweden	5 - digit	Kommuner (290)	
Switzerland	4 - digit	Gemeinden (2225)	
Taiwan	3 - digit	Township Districts (368)	
Turkey	5 - digit	Iller (81)	Ilceler (970)
United Kingdom	4 - digit	Districts (391)	

Procedure: How to Edit Geographic Configuration Settings

By default, the server is configured for unified geographic roles. This configuration is controlled by the following settings.

1. Navigate to the Workspace page, click *Settings*, point to *Geo Services*, and click *Edit settings*.

The Change Settings for GEO services page opens, as shown in the following image.

Change Settings for GEO services

Save settings in Select

▼ Customize data type mappings

? GEOROLES-AUTO Activate GEOGRAPHIC_ROLE assignment

▼ Request transformation settings

? GEO_UNIFIED_ROLE Compatibility with the previously shipped releases

? GEO_MAP_PROVIDER Map providers (Requires GEO_UNIFIED_ROLE set to ON)

? GEO_ALIASING Enables aliasing to map administrative entities

? GEO_ALIASING_APP Application for user added csv files with aliases. Default: DEFAULT

? GEO_ALIASING_FIPS Enables FIPS aliasing to map administrative entities

? GEO_COUNTRY Default value for automatically added country virtual field. Default: United States

? GEOROLE_2_MBRLEV... Enable customized per country mapping of geo roles to MB adminis

2. The following setting controls whether geographic role assignments are automatically assigned using field name analysis.

GEOROLES-AUTO. The default value is On. For example, if an alphanumeric field has the name COUNTRY, it will be assigned the geographic role Country.

3. The following settings control unified geographic roles.

- GEO_UNIFIED_ROLE.** This compatibility setting must be turned on (the default) to activate unified geographic roles. Turn it off to use geographic roles from prior releases. On indicates that the new shorter set of geographic roles that combines subsets of previously used roles will be used. Off will use geographic roles from previously shipped releases.
- GEO_MAP_PROVIDER.** Assigns names of the providers of geographic maps. The list of names should be separated with slashes to be used by transformation code for mapping. Default is *Esri and WFRS MBR*. The currently supported set of providers includes WebFOCUS Reporting Server (WFRS) (geographic boundaries distributed in the Reporting Server) and ESRI.

❑ **GEO_ALIASING.** On enables aliasing to map administrative entities. Aliasing is a mechanism to support alternative names/spellings to the administrative names used as keys to find the corresponding geometry. Caution, incorrect results will be reported on a map in the case where the column data contains variations of names/spellings for the same administrative entity (resolved to the same key). Alias names are stored in .csv files. Shipped aliasing files are located in the geomaps sub-directory of the etc folder under EDHOME. The naming convention is `geo_srv_dbl_geo_role`. Each data file has an associated synonym with the same name. Aliasing is currently supported for the four geo roles COUNTRY, STATE, COUNTY, and CITY. The STATE aliasing data file includes the valid country name for each state. COUNTY and CITY files include valid country and state names. *Valid* means the actual key value used to fetch a geometry. The default value is On.

❑ **GEO_ALIASING_APP.** Sets an application name for user-added .csv files with aliases. DEFAULT means no user files. User-added alias data files are supported for the four geo roles COUNTRY, STATE, COUNTY, and CITY (see the description for the GEO_ALIASING setting). The default value is DEFAULT, which uses the server aliasing files.

The best practice is to copy the desired .csv file from the `_edahome/etc/geomaps` folder to the application folder named in this setting, and edit it, changing aliases or adding records with new ones. There are four focexecs named `geo_srv_mapkey_<geo_role>` that have mandatory parameter ISO2 country name. The following is a request example that reports city names for South Africa sorted by state/province name.

```
EX geo_srv_mapkey_city ISO2='ZA'
```

Following is an example of user-created records (based on the obtained report):

```
"South Africa", "Gauteng", "Johannesburg", "City of Johannesburg"
"South Africa", "Western Cape", "Paarl", "Drakenstein"
```

❑ **GEO_ALIASING_FIPS.** On enables FIPS aliasing to US map administrative entities. This setting requires GEO_ALIASING to be on. The data file with the United State FIPS aliases is processed as an extension to the main aliases data file The default value is off.

❑ **GEO_COUNTRY.** Assigns a default value for an automatically added DEFINE field with GEOGRAPHIC_ROLE 'Country'. This mechanism is in effect during the create metadata process when the setting GEORoles-AUTO is ON and a column with geo role COUNTRY is not detected. The generated DEFINE field will be used to create geo hierarchies required in mapping. The default value is 'United States'.


- GEOROLE_2_MBRLEVEL.** *Enable customized per country mapping* enables customized per country mapping of the geo roles (STATE, COUNTY, and CITY) to the MBR administrative levels. This is the default value. *Disable customized per country mapping* sets uniform mapping of geo roles (STATE, COUNTY, and CITY) for all countries to the MBR administrative levels (1, 3, and 5 respectively). There is a `geo_role2mbrlev` focexec to list countries with customized admin levels. The report does not depend on this setting.

4. Click Save.

Editing the Geographic Configuration

The GEO configuration editor provides a tool for editing or adding properties for geographic roles, basemaps, reference layers, and demographic layers. In addition, it enables you to add maps and shape files to the configuration.

Go to the *Workspace* page. Click *Settings*, point to *Geo Services*, then click *Edit Configuration*. The *GEO configuration editor* opens displaying the configured geographic roles, as shown in the following image.



You can select the following objects from the Object drop-down list.

- Role.
- Basemap.
- ContextLayer.

You can edit the properties for a basemap or context layer (reference layer or demographic layer), or add a new one. You cannot edit the properties of a standard geographic role. To add a customized geographic role, you first add an Esri map or a shapefile (WebFOCUS Reporting Server map) and assign the geographic role to the map or shapefile.

Reference: Editing the List of Geographic Roles

The following standard unified geographic roles are configured by default and cannot be changed. These geographic roles create a hierarchy that can be used to drill down or up between levels of administration in maps, reports, or charts.

- CONTINENT.
- COUNTRY.
- STATE.
- COUNTY.
- CITY.
- POSTAL CODE.

To add a geographic role to the configuration, you can add a new Esri map or a shapefile hosted by the Server and associate a geographic role with the new map. You can also implement NUTS geographic roles support.

Note: Some users may not be able to run maps with custom geographic roles when the default reporting server security configuration is in place. To allow a user to run content with custom geographic roles, on the Access Control page of the Reporting Server browser interface, right-click the role that the user privileges are based on and click *Directory/File Privileges*, then enable both the Execute and List privileges for `geo_services_user.xml`.

Configuring Properties of Geographic Roles

The following describes columns for geographic roles in the configuration editor.

`name`

Is the unique name of the geographic role. It cannot have spaces, but it can have underscores (_).

Next to the name is an indicator of whether the role is a standard role or a customized role.

`title`

Is the description of the geographic role that is displayed in reports and in drop-down lists in the WebFOCUS tools.

`returned_geometry`

Is the type of geographic data returned from the map service for rendering on the map.

Valid values include:

- GEOMETRY_AREA
- GEOMETRY_POINT
- GEOMETRY_LINE

When you have configured the properties, click *OK* to return to the Geo Configuration Editor.

The new role will display on the list of roles, Click *Save* to save it in the geographic configuration.

Adding a New Role for an Esri Map

To add an Esri geographic role, click *Add ESRI map* on the Geographic Configuration Editor.

The *Create a new ESRI map role* dialog box opens, as shown in the following image.

Create a new ESRI map role

Associated mapping service properties

name:

title:

returned_geometry: GEOMETRY_AREA

provider: ESRI

authorization: none

url:

Verify

Service parameters:

position	GEO role	parameter name
1		

OK Cancel

Configure the following map service properties.

name

Is a name for the geographic role.

title

Is a title to display in the WebFOCUS tools.

returned_geography

Select the type of geometry that is returned from the map service for this role. Valid values are:

- GEOMETRY_AREA.** Returns JSON polygon definitions.
- GEOMETRY_LINE.** Returns JSON line definitions.
- GEOMETRY_POINT.** Returns a JSON point.

url

Is the URL to the map service that provides the geographic data.

Click *Verify* after entering the URL to verify that the map service is available by going to the specified URL.

Service Parameters

Add as parameters any additional geographic roles needed to identify the exact location of the new role. For example, a city name needs state and country parameters.

Example: Adding the WebFOCUS Regions Geographic Role

The following properties add the WebFOCUS Regions role to the configuration.

Associated mapping service properties

name: REGION

title: WF_REGIONS

returned_geometry: GEOMETRY_AREA

provider: ESRI

authorization: none

url: http://services7.arcgis.com/L95Wwv90jRQ0tjAs/ArcGIS/rest/services/wfretail_sub_regions/FeatureSe

Verify

Service parameters: Add

position	GEO role	parameter name
1	REGION	SUB_REGION
2		

Note: The *parameter name* corresponds to the field name in the FeatureLayer referenced in the following URL:

http://services7.arcgis.com/L95Wwv90jRQ0tjAs/ArcGIS/rest/services/wfretail_sub_regions/FeatureServer/0

Click *OK* when you have finished configuring the properties.

The new role is added to the configuration as a customized role, as shown in the following image.

GEO configuration editor

Object:

(Edit to check other parameters of a role) ? Help + Add ESRI map + Add WFRS map + Add NUTS support + Add extended postal codes

name		title	returned_geometry	provider
REGION	customized	WF_REGIONS	GEOMETRY_AREA	ESRI
CONTINENT	standard	Continent	GEOMETRY_AREA	ESRI
COUNTRY	standard	Country	GEOMETRY_AREA	ESRI
STATE	standard	State	GEOMETRY_AREA	WFRS/ESRI
COUNTY	standard	County	GEOMETRY_AREA	WFRS/ESRI
CITY	standard	City	GEOMETRY_POINT	WFRS/ESRI
POSTAL-CODE	standard	Postal code	GEOMETRY_AREA	WFRS/ESRI

Click *Save* to save this role to the configuration.

The following request uses the WebFOCUS Regions geographic role in a map request.

```


DEFINE FILE WF_RETAIL_LITE

REGION/A50 (GEOGRAPHIC_ROLE=REGION) = BUSINESS_SUB_REGION;
END

GRAPH FILE WF_RETAIL_LITE
SUM COGS_US
BY REGION
WHERE COUNTRY_NAME EQ 'United States'
ON GRAPH PCHOLD FORMAT JSCHART
ON GRAPH SET LOOKGRAPH CHOROPLETH
ON GRAPH SET STYLE *
TYPE=REPORT, CHART-LOOK=com.esri.map, $
TYPE=DATA, COLUMN=N2, BUCKET=color, $
*GRAPH_JS_FINAL
"extensions": {
"com.esri.map": {
"overlayLayers":
[
{
"ibiDataLayer": {
"map-metadata": {
"map_by_field": "REGION"
}
}
}
],
"baseMapInfo":
{
"customBaseMaps":
[
{
"ibiBaseLayer": "gray"
}
]
}
}
}
*END
ENDSTYLE
END

```

The output is shown in the following image.



Adding a New Role for a Server-Hosted Map

A server-hosted map is based on a shapefile. You must upload the shapefile (.dbf) to an application folder accessible to the server. The server will transform it to ibijson format.

An ESRI shape file is actually a collection of at least four files:

- .dbf file.** The .dbf file is a standard database file used to store attribute data and object IDs. A .dbf file is mandatory for shape files.
- .shp file.** The .shp file is a mandatory Esri file that gives features their geometry. Every shapefile has its own .shp file that represents spatial vector data.
- .shx file.** The .shx file is a mandatory Esri shape index position file. This type of file is used to search forward and backward.
- .prj file.** The .prj file is an optional file that contains the metadata associated with the shapefiles coordinate and projection system.

All files must have exactly the same name and to be located in the same directory. If they are not, the shapefile conversion will fail. Additionally, the names must not contain spaces.

When there are several possible keys associated with a geometry, a drop-down list of detected key names will be displayed. Select any one of these fields. No selection required when there is a single geometry key.

The shapefile should only be in the GCS_WGS_1984 - World Geodetic System 1984 (decimal degrees) coordinate system.

To add a geographic role for a Server-hosted map, click *Add WFRS map* on the Geographic Configuration Editor.

The *Add WFRS hosted map* dialog box opens, as shown in the following image.

Configure the following properties.

`role name`

Is a name for the geographic role.

`Geometry type`

Select either POLYGON or POINT from the drop-down list.

`Esri shape`

Enter the name of the application directory where the shapefile resides, or click the ellipsis (...) to navigate to the application directory. Then select the .dbf file for the role.

`Load to app`

Enter the name of the application directory where you want to place the ibijson file, or click the ellipsis (...) to navigate to the application directory.

`Quantization type`

Quantization is the process of transforming a large set of input values to a smaller set of values. When transforming the shapefile, the server will quantize points that are too close together in order to optimize map rendering performance. Two methods are available for quantization, LINEAR or GRID. The default is LINEAR.

`Quantization_X`

Is the threshold value for the x-axis.

Quantization_Y

Is the threshold value for the y-axis.

If the map has multiple keys, a drop-down list displays so that you can select one.


Click *OK* when you have finished configuring the properties.

The new role is added to the configuration as a customized role, as shown in the following image.



Click *Save* to save this role to the configuration.

You can test the role by right-clicking the role in the configuration editor and clicking *Test*. A sample map will be generated, as shown in the following image.



Adding NUTS Support

Nomenclature of territorial units for statistics (NUTS) are geographic roles specific to the European Union.

To add NUTS geographic roles to the configuration, click *Add NUTS support* on the Geographic Configuration Editor.

The NUTS geographic roles are added, as shown in the following image.

GEO configuration editor

Object:

(Edit to check other parameters of a role) ? Help + Add ESRI map + Add WFRS map + Add NUTS support + Add extended postal codes

name		title	returned_geometry	provider
NUTS0	customized	Country (NUTS level 0)	GEOMETRY_AREA	ESRI
NUTS0_CC	customized	Country (NUTS level 0)	GEOMETRY_AREA	ESRI
NUTS1	customized	Region (NUTS level 1)	GEOMETRY_AREA	ESRI
NUTS1_CC	customized	Region (NUTS level 1)	GEOMETRY_AREA	ESRI
NUTS2	customized	Province (NUTS level 2)	GEOMETRY_AREA	ESRI
NUTS2_CC	customized	Province (NUTS level 2)	GEOMETRY_AREA	ESRI
NUTS3	customized	District (NUTS level 3)	GEOMETRY_AREA	ESRI
NUTS3_CC	customized	District (NUTS level 3)	GEOMETRY_AREA	ESRI
CONTINENT	standard	Continent	GEOMETRY_AREA	ESRI
COUNTRY	standard	Country	GEOMETRY_AREA	ESRI
STATE	standard	State	GEOMETRY_AREA	WFRS/ESRI
COUNTY	standard	County	GEOMETRY_AREA	WFRS/ESRI
CITY	standard	City	GEOMETRY_POINT	WFRS/ESRI
POSTAL-CODE	standard	Postal code	GEOMETRY_AREA	WFRS/ESRI

Click **Save** to save these roles to the configuration.

Adding Support for Extended Postal Codes

Click *Add extended postal codes* to add support for Level 1 and Level 2 postal codes used in certain countries.


Customizing Vocabulary Rules

For each geographic role, a set of vocabulary rules define how to recognize when a field name should automatically be assigned to that role. If you right-click a role, you can click *Customize vocabulary* from the shortcut menu.

Elements in a rule are connected by the Boolean logic operation OR (only one needs to be satisfied). Each vocabulary element contains words enclosed with special characters. Words in the rule element are connected by the Boolean logic operation AND (all need to be satisfied).

A word may be prefixed and/or suffixed with the percent character (%), which is a placeholder for any sequence of characters. If an element contains more than one word, each word has to be prefixed by the character plus (+) or minus (-). Plus indicates that the word must be found in the column name. Minus indicates that word must not be found in the column name.

For example, the following are the vocabulary rules for the role COUNTRY.



To add another rule, click *Add optional*.

When you are finished, click *OK*.

Click *Save* to save these rules to the configuration.

Reference: Customizing the List of Basemaps

You can edit an existing basemap definition or add a custom basemap.

Basemaps that are not vector maps are pre-rendered image tiles. Vector basemaps are delivered from ArcGIS Online as vector tiles that are rendered on the client based on a style file that is delivered with the vector tiles. This provides customization options not available with pre-rendered image tiles.

A default set of vector basemaps is configured in the GEO configuration editor in the Server. However, you can configure additional vector basemaps as needed. You can find the basemap URLs and Style URLs at <http://www.arcgis.com>.

Using Standard Basemaps

To add a new basemap or customize an existing basemap, select *Basemap* from the Object drop-down list.

The following image shows the GEO configuration editor with the Basemap object selected, showing the list of standard basemaps.

GEO configuration editor

Object:

+ Add						
name		icon	title	type	url	addon_json
streets	standard	qb/streets_map_108x72.png	World Street Map	tiled		
satellite	standard	qb/imagery_map_108x72.png	World Imagery	tiled		
terrain	standard	qb/terrain_map_108x72.png	Terrain with Labels	tiled		
gray	standard	qb/gray_map_108x72.png	Light Gray Canvas Map	tiled		
dark-gray	standard	qb/darkgray_map_108x72.png	Dark Gray Canvas Map	tiled		
oceans	standard	qb/oceans_map_108x72.png	Oceans Map	tiled		
national-geographic	standard	qb/national_geographic_map_108x72.png	National Geographic World Map	tiled		
osm	standard	qb/open_street_map_108x72.png	Open Street Map	tiled		
hybrid	standard	qb/hybrid_map_108x72.png	World Imagery with Labels	tiled		
topo	standard	qb/topo_map_108x72.png	World Topographic Map	tiled		
gray-vector	standard	qb/light_gray_vector.jpg	Gray Vector	vector		
streets-night-vector	standard	qb/streets_night_map.jpg	Streets Night Vector	vector		
streets-navigation-vector	standard	qb/navigation_day_map_vector.jpg	Streets Navigation Vector	vector		
topo-vector	standard	qb/topographic_relief_vector.jpg	Topographic Vector	vector		
dark-gray-vector	standard	qb/dark_gray_vector.jpg	Dark Gray Vector	vector		
streets-vector	standard	qb/streets_map_vector.jpg	Streets Vector	vector		
streets-relief-vector	standard	qb/street_relief_vector.jpg	Streets Relief Vector	vector		
None	standard	qb/none_map_108x72.png	BasemapNone	tiled		

Configuring Basemap Properties

The following is a description of the properties used for basemap configuration.

name

Is the name of the basemap.

Next to the name is an indicator of whether the basemap is a standard basemap or a customized basemap.

icon

Is the name of the thumbnail for the basemap (for a standard basemap) or the URL to the thumbnail (for a customized basemap) that will appear on the Basemap drop-down list in the WebFOCUS tools or the Change Basemap map widget.

title

Is a title to display on the Basemap drop-down list in the WebFOCUS tools or the Change Basemap map widget.

url

Is the URL to the map service that provides the basemap, for a customized basemap. The map service URL can be copied from the URL field on the page showing a custom basemap. For a standard basemap, the URL is already stored in the server geographic configuration file and is not displayed.

type

Valid values are tiled and vector.

addon_json

Specifies additional JSON properties for rendering the map. For a custom vector basemap, enter the basemap map service URLs and Style URLs needed for rendering the map. The map can consist of an array of layers, each with a URL and its style URL. The style URL for each layer can be found by clicking *View Style* on the page showing the map layer. The map service URL for each layer can be copied from the URL field on the page showing the map layer. Enter the information in the following format.

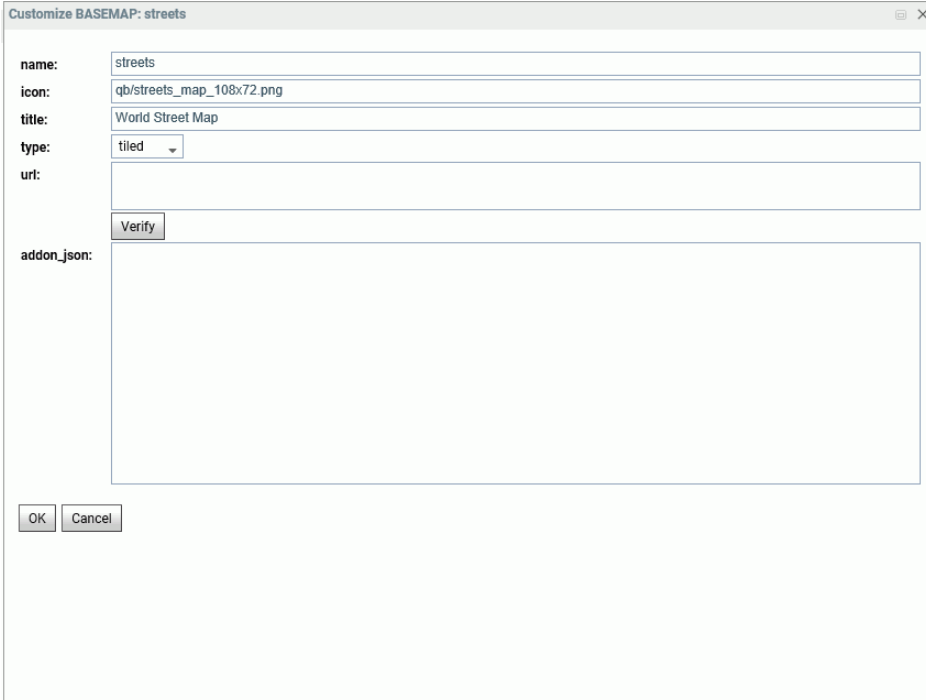
```
{ "baseMapLayers":  
  [  
    { "url": "url_to_map_service", "styleUrl": "url_to_style" },  
    { "url": "url_to_map_service", "styleUrl": "url_to_style" },  
    ...  
  ]  
}
```

Note: If the vector basemap has multiple layers, its page on ArcGIS online will have links to each layer.

Editing the Properties of a Standard Basemap

To customize the properties of an existing basemap, right-click the basemap line and click *Customize BASEMAP*.

The Customize Basemap dialog box opens, as shown in the following image.



Customize BASEMAP: streets

name: streets

icon: qb/streets_map_108x72.png

title: World Street Map

type: tiled

url:

Verify

addon_json:

OK Cancel

Edit the properties you want to change. If you change the URL, you can click *Verify* to make sure the map service is valid and accessible.

When you are finished, click *OK*, then click *Save* on the GEO configuration editor Basemaps page.

Adding a Custom Basemap

To add a new basemap to the configuration, click *Add*.

The *Create a BASEMAP* dialog box opens, as shown in the following image.

The image shows a dialog box titled "Create a BASEMAP". It has a standard window title bar with a close button. The dialog contains the following fields and controls:

- name:** A text input field.
- icon:** A text input field.
- title:** A text input field.
- type:** A dropdown menu currently showing "tiled".
- url:** A text input field.
- addon_json:** A large text area for additional configuration.
- Verify:** A button located below the "url:" field.
- OK** and **Cancel:** Buttons at the bottom left of the dialog.

Enter a name for the basemap, a URL to the thumbnail, a title to display, and the URL to the map service that provides the basemap, and click *Verify*.

When you have configured the properties, click *OK*, then click *Save* on the GEO configuration editor Basemap page.

Example: **Configuring the Nova Vector Basemap**

The following steps are needed to determine the properties associated with the Nova vector basemap.

1. Navigate to the following URL for the Nova basemap.

<https://www.arcgis.com/home/item.html?id=75f4dfdff19e445395653121a95a85db>.

The Nova basemap page is shown in the following image.

The screenshot displays the 'Overview' page for the 'Nova' basemap. At the top, there's a navigation bar with 'Overview' selected. Below it, a small thumbnail shows a map of the Northeastern United States. To the right of the thumbnail, text reads: 'This (v2) vector tile layer provides a detailed basemap for the world featuring a dark background with glowing blue symbology inspired by the ArcGIS.com splash screen.' Below this, it says 'Tile Layer by Esri', 'Created: Nov 6, 2017', 'Updated: Jun 27, 2018', and 'View Count: 136'. There are two tags: 'Authoritative' and 'Living Atlas'. On the right side, there are three buttons: 'Open in Map Viewer', 'Open in Scene Viewer', and 'View style'. Below the buttons is a 'Details' section with 'Source: Vector Tile Service', 'Size: 343 KB', and a star rating. There are social media icons for Facebook, Twitter, and LinkedIn. The 'Owner' section shows the Esri logo. The 'Managed by' section shows the Esri logo. The 'Tags' section lists: 'v2, basemaps, vector, creative, general availability, esri_basemap, basemap, Cindy Prostak'.

In the `addon_json` syntax, you need the `styleUrl` for each basemap URL defined in the vector basemap.

2. Click the *View Style* button to retrieve the `styleUrl`. In this case, it is the following URL.

<https://www.arcgis.com/sharing/rest/content/items/75f4dfdff19e445395653121a95a85db/resources/styles/root.json?f=pjson>

3. Add the following URL for the Esri vector map service in both the `url` field in the GEO configuration editor and in the `addon_json` syntax for this basemap. You can copy this location from the URL field on the lower right of the Nova basemap page.

https://basemaps.arcgis.com/arcgis/rest/services/World_Basemap_v2/VectorTileServer

4. The `icon` value must be a URL reference. This example uses the one referenced in the thumbnail of the webmap, which is the following URL (retrieved by right-clicking the thumbnail image at the top of the page and copying the image location).

<https://www.arcgis.com/apps/mapviewer/index.html?layers=75f4dfdff19e445395653121a95a85db>

Add these properties in the GEO configuration editor to define the Nova vector basemap, which consists of a single layer. Add the name, type, and a title, as follows.

name: Nova

icon: <https://www.arcgis.com/apps/mapviewer/index.html?layers=75f4dfdff19e445395653121a95a85db>

title: Nova

url: https://basemaps.arcgis.com/arcgis/rest/services/World_Basemap_v2/VectorTileServer

addon_json:

```
{ "baseMapLayers" :  
  [  
    { "url": "https://basemaps.arcgis.com/arcgis/rest/services/World_Basemap_v2/  
VectorTileServer",  
      "styleUrl": "http://www.arcgis.com/sharing/rest/content/items/  
75f4dfdf19e445395653121a95a85db/resources/styles/root.json"  
    }  
  ]  
}
```

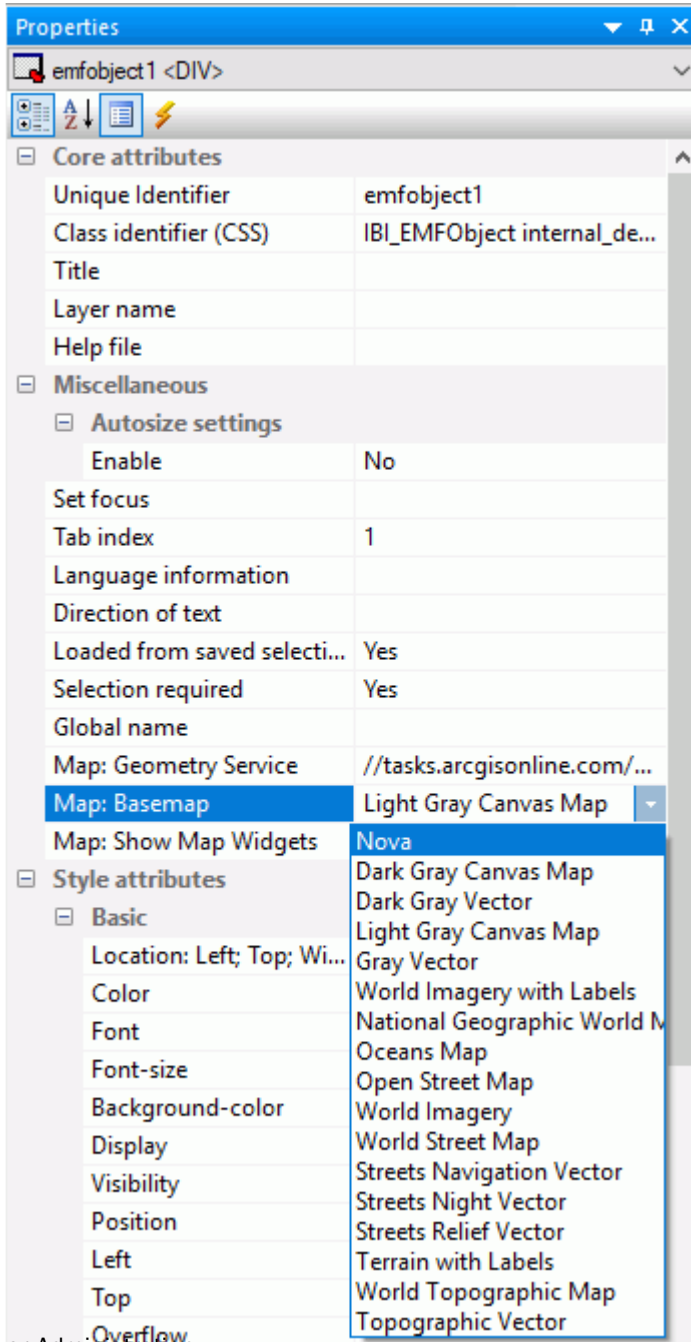
The following image shows the Create a BASEMAP page for configuring the Nova vector basemap in the Reporting Server GEO configuration editor.

The screenshot shows a window titled "Create a BASEMAP" with the following fields and content:

- name:** Nova
- icon:** <http://www.arcgis.com/sharing/rest/content/items/75f4dfdf19e445395653121a95a85db/info/thumbnail/Nova.png>
- title:** Nova
- type:** vector
- url:** https://basemaps.arcgis.com/arcgis/rest/services/World_Basemap_v2/VectorTileServer
- Verify** button
- addon_json:**

```
{ "baseMapLayers":  
  [  
    { "url": "https://basemaps.arcgis.com/arcgis/rest/services/World_Basemap_v2/VectorTileServer",  
      "styleUrl": "http://www.arcgis.com/sharing/rest/content/items/75f4dfdf19e445395653121a95a85db/reso  
" }  
  ]  
}
```
- OK** and **Cancel** buttons at the bottom.

After clicking OK, then Save, the custom vector basemap becomes available for selection in the WebFOCUS® App Studio Properties panel for the Esri map component, as shown in the following image.




The WebFOCUS App Studio Esri component in this example uses the following request, named `customer_locations.fex`.


```
TABLE FILE WF_RETAIL_LITE
PRINT
    ID_CUSTOMER
    COGS_US
    QUANTITY_SOLD
    COMPUTE LONGITUDE/D12.2 =CUSTOMER_LONGITUDE;
    COMPUTE LATITUDE/D12.2 = CUSTOMER_LATITUDE;
BY STATE_PROV_NAME
BY ID_CUSTOMER NOPRINT
WHERE COUNTRY_NAME EQ 'United States';
ON TABLE NOTOTAL
ON TABLE PCHOLD FORMAT XML
END
```

Configure the Settings panel for the Esri map to use this request for a Point of Interest layer.

Running the HTML page generates the following map, with the Nova vector basemap as the base layer.



The Nova vector basemap is also available in the Change Basemap widget, as shown in the following image.



Reference: Customizing the List of Context Layers

To add a new context layer or customize an existing context layer, select *ContextLayer* from the Object drop-down list.

The following image shows the GEO configuration editor with the ContextLayer object selected.

GEO configuration editor

Object: ContextLayer

+ Add

name	authorization	layerType	title	uri
USA_Population_Density_2012	standard	none	tile	USA Population Density 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Population_Density/MapServer
USA_Population_Density_2017	standard	silent	featurelayer	USA Population Density 2017 http://demographics7.arcgis.com/arcgis/rest/services/USA_Demographics_and_Boundaries_2017/MapServer
USA_Population_Change_1990_2000	standard	none	tile	USA Population Change 1990-2000 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_1990-2000_Population_Change/MapServer
USA_Population_Change_2000_2010	standard	none	tile	USA Population Change 2000-2010 http://server.arcgisonline.com/arcgis/rest/services/Demographics/USA_2000-2010_Population_Change/MapServer
USA_Population_Change_2010_2012	standard	none	tile	USA Population Change 2010-2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Recent_Population_Change/MapServer
USA_Projected_Population_Growth_2012_2017	standard	none	tile	USA Projected Population Growth 2012-2017 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Projected_Population_Change/MapServer
USA_Population_Growth_2017_2022	standard	silent	featurelayer	USA Population Growth 2017-2022 http://demographics7.arcgis.com/arcgis/rest/services/USA_Demographics_and_Boundaries_2017/MapServer
USA_Median_Household_Income_2012	standard	none	tile	USA Median Household Income 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Median_Household_Income/MapServer
USA_Population_Older_than_Age_64	standard	none	tile	USA Population Older than Age 64 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Percent_Over_64/MapServer
USA_Unemployment_Rate_2012	standard	none	tile	USA Unemployment Rate 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Unemployment_Rate/MapServer
USA_Unemployment_Rate_2017	standard	silent	featurelayer	USA Unemployment Rate 2017 http://demographics7.arcgis.com/arcgis/rest/services/USA_Demographics_and_Boundaries_2017/MapServer
USA_Median_Home_Value_2012	standard	none	tile	USA Median Home Value 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Median_Home_Value/MapServer
USA_Population_Younger_than_Age_18	standard	none	tile	USA Population Younger than Age 18 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Percent_Under_18/MapServer
USA_Average_Household_Size_2012	standard	none	tile	USA Average Household Size 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Average_Household_Size/MapServer
USA_Median_Net_Worth_2012	standard	none	tile	USA Median Net Worth 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Median_Net_Worth/MapServer
USA_Owner_Occupied_Housing_2012	standard	none	tile	USA Owner Occupied Housing 2012 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Owner_Occupied_Housing/MapServer
USA_Labor_Force_Participation_Rate_2010	standard	none	tile	USA Labor Force Participation Rate 2010 http://server.arcgisonline.com/ArcGIS/rest/services/Demographics/USA_Labor_Force_Participation_Rate/MapServer
USA_Population_By_Sex_2012	standard	none	tile	USA Population By Sex 2012 http://server.arcgisonline.com/arcgis/rest/services/Demographics/USA_Population_by_Sex/MapServer
USA_Median_Age_2017	standard	silent	featurelayer	USA Median Age 2017 http://demographics7.arcgis.com/arcgis/rest/services/USA_Demographics_and_Boundaries_2017/MapServer

Configuring Context Layer Properties

Following is a description of the properties used for context layer configuration.

name

Is the name of the context layer.

Next to the name is an indicator of whether the context layer is a standard context layer or a customized context layer.

authorization

Is the type of authentication needed to access this context layer. Valid values are:

silent. Credentials for your ArcGIS application are provided in the connection string of the Adapter for Esri ArcGIS.

Note: For instructions for configuring the Adapter for Esri ArcGIS, see the *TIBCO WebFOCUS® Adapter Administration* manual.

none. No authorization is needed.

named. User credentials are provided in the connection string of the Adapter for Esri.

on premises. User credentials for a locally hosted ArcGIS server are provided in the connection string of the Adapter for Esri.

`layer type`

Is the type of context layer. For a cached layer, the layer type is `tile`. For a layer that is rendered dynamically, the layer type is `featurelayer`.

`title`

Is a title to display on the demographic layer drop-down list in the WebFOCUS tools.

`addon_json`

Specifies additional JSON properties that can be used to control the rendering of the context layer. For example, `smartMapping` properties define the border styles within the context layer.

Depending on the type of context layer, the `addon_json` properties may resemble the following:

```
"layerObjectType" : "esri/layers/ArcGISDynamicMapServiceLayer",
  "smartMapping": { "webMapInfo": {
    "queryString" : "id:id",
    "itemDataUrl" : "url",
    "layerTypeEx" : "layerType"
  }}
}
```

where:

`esri/layers/ArcGISDynamicMapServiceLayer`

Identifies the context layer as a dynamic map service.

id

Is the ID of the context layer on Esri ArcGIS. This can be found in the URL for the layer on the Esri ArcGIS site, or as the value of the `"serviceltemId"` property in the JSON properties accessible from the service URL of the layer.

If there is no layer ID, then this property can be omitted.

url

Is the map service URL for the layer. This should match the value that you provide for the `uri` field.

layerType

Specifies the type of context layer. It can be one of the following values:

- group.** Identifies the context layer as a group layer composed of multiple feature layers. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-GroupLayer.html>.

- ❑ **map-image.** Identifies the context layer as a map image layer, which can contain multiple sub-layers rendered as images rather than features. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-MapImageLayer.html>.
- ❑ **feature.** Identifies the context layer as a feature layer, which is a single layer created from a map service or feature service. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-FeatureLayer.html>.
- ❑ **imagery-tile.** Identifies the context layer as an imagery tile layer, created from a tiled image service. For more information see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-ImageryTileLayer.html>.
- ❑ **imagery.** Identifies the context layer as an imagery layer, created from an image service. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-ImageryLayer.html>.
- ❑ **elevation.** Identifies the context layer as an elevation layer used in a 3D SceneView. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-ElevationLayer.html>.
- ❑ **streaming.** Identifies the context layer as a streaming layer. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-StreamLayer.html>
- ❑ **tile.** Identifies the context layer as a tile layer, created from a cached map service. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-layers-TileLayer.html>.
- ❑ **web-map.** Identifies a complete, two-dimensional WebMap. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-WebMap.html>.
- ❑ **web-scene.** Identifies a complete, 3D web scene. For more information, see <https://developers.arcgis.com/javascript/latest/api-reference/esri-WebScene.html>.

For example, the add-on JSON for a map image layer may resemble the following:

```
"layerObjectType" : "esri/layers/ArcGISTiledMapServiceLayer",
"smartMapping": { "webMapInfo": {
  "queryString" : "id:1234567890abcdefghij1234567890ab",
  "itemDataUrl" :
  "https://services5.arcgis.com/1234567890abcdef/arcgis/rest/services/
Layer_Name/FeatureServer",
  "layerTypeEx" : "map-image"
}
}
```

`uri`

Is the URL to the map service that provides the context layer.

Customizing the Properties of an Existing Context Layer

To customize the properties of an existing context layer, click the down arrow next to a context layer name or right-click the context layer line and click *Customize context layer*.

The Create a Custom CONTEXTLAYER dialog box opens. The completed dialog box may resemble the one shown in the following image.

The screenshot shows a dialog box titled "Customize CONTEXTLAYER: USA_Population_Density_2017". It has the following fields and values:

- name:** USA_Population_Density_2017
- authorization:** silent
- layerType:** featurelayer
- title:** USA Population Density 2017
- addon_json:** {"layerObjectType": "esri/layers/ArcGISDynamicMapServiceLayer", "smartMapping": { "webMapInfo": {
- uri:** http://demographics7.arcgis.com/arcgis/rest/services/USA_Demographics_and_Boundaries_2017/MapServer

At the bottom of the dialog box, there are three buttons: "Verify", "OK", and "Cancel".

Edit the properties you want to change. If you change the URI, you can click *Verify* to make sure the map service is valid and accessible.

When you are finished, click *OK*, then click *Save* on the GEO configuration editor Context Layers page.

Adding a New Context Layer

To add a new context layer to the configuration, click *Add*.

The **Create a Customized CONTEXTLAYER** dialog box opens, as shown in the following image.

The image shows a dialog box titled "Create a Customized CONTEXTLAYER". It has a standard window title bar with a close button. The dialog contains the following fields and controls:

- name:** A text input field.
- authorization:** A dropdown menu with "none" selected.
- layerType:** A dropdown menu with "tile" selected.
- title:** A text input field.
- addon_json:** A large text area for entering JSON data.
- uri:** A text input field.
- Verify:** A button located below the uri field.
- OK** and **Cancel:** Buttons at the bottom left.

Enter a name for the context layer, the authorization type, a layer type, a title to display, any additional JSON needed for rendering the context layer, and the URI to the map service that provides the context layer, and click *Verify*.

When you have configured the properties, click *OK*, then click *Save* on the GEO configuration editor Context Layer page.

Migrate Settings

The Migrate Settings page allows you to migrate a server from a previous release.

It is recommended that you migrate configuration information from previous releases after you verify the proper basic installation of the new release.

To migrate from a previous release:

1. Go to the *Workspace* page.
2. Click *Settings*, then click *Migrate*.

The Migrate page opens.

The screenshot shows a web form with a text input field. The label for the field is "? Enter EDACONF directory to migrate from". The text entered in the field is "C:\ibisrv77\wfsold". To the right of the input field is an ellipsis "...". Below the input field is a button labeled "Continue".

3. Type the full path of the configuration instance directory (EDACONF) or click the ellipsis (...) and navigate to it.
4. Click *Continue*.

Monitoring TIBCO WebFOCUS Reporting Server Activity

You can monitor every server type and change its operation parameters through the Reporting Server browser interface. Those changes can affect the behavior and performance of the server.

This topic describes how to monitor the following:

- Data Services
- Java Services
- Special Services and Listeners

It also describes how to perform certain administrative tasks on a server.

Monitoring TIBCO WebFOCUS Reporting Server Status

You can view the system and system memory status from the Status page.

1. To access the Status page, go to the *Workspace* page.
2. Click *Monitor*, then click *Status*.

The Status page opens, as shown in the following image.

^ System	
Workspace manager process id	5312
Started on	2018/03/28 13:06:19
EDASHARE	C:\ibi\srv77\wfs
EDATEMP	C:\ibi\srv77\wfs\edatemp
EDATEMP available disk space (KB)	339522764
CPU Usage	16%

^ System Memory Status	
RAM (KB)	8272332
RAM free (KB)	4346004
paging file (KB)	16660940
free paging file (KB)	11186872
virtual memory (KB)	137438953344
free virtual memory (KB)	137438784340

Monitoring Data Services

From the Reporting Server browser interface, you can:

- Display the Data Services Agents page, and customize the page to include only those statistics that you are interested in. The Data Services Agents page displays statistics for the current list of agents monitored by the Workspace Manager.
- Display the Data Service Agents page for a selected service. You can customize the page to include only those statistics that you are interested in.
- Display statistics, including system-extended statistics, for an individual Data Service Agent.
- Display a status report for all Data Services. You can customize the report to include only those statistics that you are interested in.
- Display aggregated statistics for all Data Services.

- Display aggregated statistics for a selected Data Service.
- Display connection statistics.

As you monitor Data Service activity, you can also perform the following administrative tasks from the Reporting Server browser interface:

- Reset aggregated statistics for Data Services and start collecting them from the beginning.
- Kill one or more Data Service Agents.
- Start one or more new Data Service Agents.
- Display an error message for a Data Service Agent.

Displaying and Customizing a Data Service Agents Page

This topic describes how Server Administrators can display the Data Service Agents page for all Data Services or for an individual Data Service. Server Administrators access the page from the Workspace ribbon, using the *Monitor* menu.

If you are a user other than a Server Administrator, you can access the same page and the same statistics from My Console. With the Manage My Agents option, you can monitor and manage your own agents, that is, the agents that match your user ID. If granted the applicable privileges, you can also monitor, or monitor and manage, the agents of the other users in your group.

To display and customize a Data Service Agents Page:

1. Go to the *Workspace* page.
2. Click *Monitor*, then click *Data Service Agents*, or right-click *Data Services* or a type of agent on the resources pane, and click *Agents*.

The Data Services Agents page opens.

The table in [Data Service Agents Performance Parameters](#) on page 663 describes all the statistics that you can display on the page.

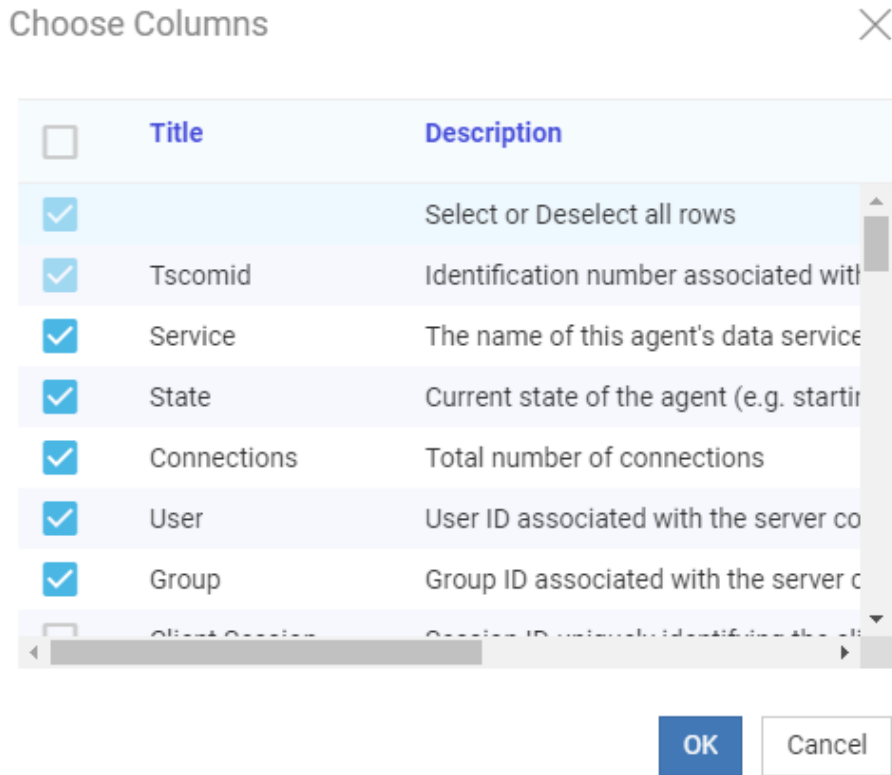
3. To choose specific statistics that you want to display on the report, click *Choose Columns* on the *View* menu (the down arrow on the menu bar).

The page that opens lists the available statistics. A description of each statistic is provided in the Description column, to help you tailor the page for your needs.

Select the check box for each statistic that you want to display on the page, and de-select the check box for each statistic that you do not want to display.

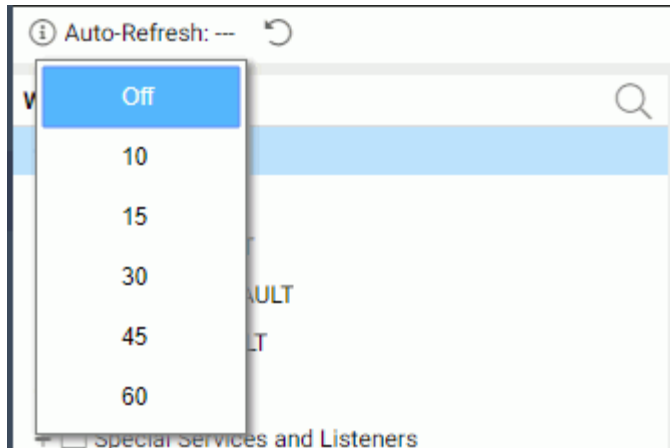
Click *OK* when you have made your selections.

The following image shows the Choose Columns dialog box.



4. To filter the values that you want to display on the page, click the relevant check boxes. You are returned to the Data Service Agents page with your filter activated.
5. To sort the list based on the order of values in a column, click the column heading. The sort toggles between ascending order (indicated by an up arrow), descending order (indicated by a down arrow), and original order (no arrow displays).
6. Click the Refresh button (next to Auto-Refresh) on the menu bar above the resources tree) to refresh the statistics. On top of the resources tree, there is also a refresh icon.

You can also set automatic refresh options by clicking Auto-Refresh on top of the resources tree, as shown in the following image.



Select a refresh interval.

Procedure: **How to Display and Customize a Data Service Agents Page Associated With a Selected Service**

Since agents are associated with services, you may wish to display current service statistics.

1. Go to the *Workspace* page.
2. In the resources pane, expand the *Data Services* folder.
3. Right-click the name of the service on which you want to report, for example, *WC_DEFAULT*.
4. From the menu, click *Agents*.

The Data Service Agents page associated with the selected service is displayed on the right. The table in [Data Service Agents Performance Parameters](#) on page 663 describes all the statistics that are provided on the page.

5. You can customize and filter the report, as described in [Displaying and Customizing a Data Service Agents Page](#) on page 661.

Reference: **Data Service Agents Performance Parameters**

The following statistics can appear on a Data Service Agents page. Each agent is represented as a separate row on the report.

General

Tscomid

Identification number associated with the Data Service Agent.

Service

Name of the data service for the agent.

State

Current state of the agent. Possible values are in use, DBMS call, idle, aborted, crashed, stopping, or starting.

A session connects and is *in use* until the moment that it disconnects. This includes all time spent between a session suspension and resumption when the process is not using the CPU but still has resources allocated for at least one session.

DBMS call indicates that the agent is waiting for a database to return data.

Idle means that the agent is not connected. It has no sessions at all, whether active or suspended.

Aborted and crashed are abnormal states resulting from a fatal software error detected by the program (aborted) or by the operating system (crashed). The agent process is no longer running, and these states are provided for diagnostic purposes. The Server Administrator can clear (using the kill option) such agents after the problem is investigated.

Stopping and starting are normal transitory states, which are self-explanatory.

Connections

Shows the number of connections that have been used for each agent. You can compare this number to the maximum number set for each service. The maximum number of agents allowed for each service can be changed by right-clicking the service and clicking *Properties*.

Client Information

User

User ID associated with the server connection.

Group

Group ID associated with the server connection.

Client Session

Session ID uniquely identifies the client session.

Client Address

For some protocols, the network address of the connected client.

Code Page

Character code page used by the client connected to the session.

Defer ID

Unique identifier assigned to the submitted deferred request.

Scheduler Job ID

Unique identifier assigned to the last submitted scheduled or CMASAP request.

FOCUS Information

Query Time

Last time that a request to the agent was made. This value is used to calculate the time that an agent is idle, in cases in which an *idle agent limit* was set.

Command

First eight characters of the last instruction executed by the data access agent.

Procedure

Displays the procedure currently being executed or completed by each agent. For example, when a procedure, which is a part of a job, executes another procedure, the name of this procedure is shown in a Procedure column.

Top Level Procedure

Displays calling procedures or calling flows executed or submitted by agents.

Note: This column is not shown by default on the Data Services Agents page. To display it, click *Choose Columns*, select its check box on the Select the columns that will appear on the Agents page, and click *OK*.

Master File

Name of the last Master File that was processed by the agent. The displayed name can be up to 64 characters.

Executed Class

Last Java class executed. This value includes explicitly called Java classes (such as CALLJAVA), and implicitly called classes (such as GRAPH FILE).

Response Time

Time from the moment that the server receives the request (regardless of the requesting tool, that is, Reporting Server browser interface, Data Migrator, or WebFOCUS) until the request completes and the response is sent to the requester.

Server Time

Response time minus DBMS processing time.

Error Number

Last &RETCODE that is set upon procedure (FOCEXEC) completion. Zero (0) is normal. Any other value indicates an abnormal completion.

Focus I/O

Number of FOCUS I/O operations performed by the agent.

Foccache Ticket

Unique identifier assigned to the last request that used the foccache directory.

Number of Transactions

Number of transactions or HLI commands performed by the agent.

Adapter Information

DBMS I/O

Number of external I/O operations performed by the agent.

DBMS Time

Amount of DBMS processing time in seconds.

Change Data Capture Information

Source DBMS Adapter

Suffix of the listening adapter.

Source Table Name

Data source that is listened on.

Status

Current status of the adapter.

Processed Transactions

Number of processed LUWs (files).

Polling Interval

Interval at which polling takes place.

Timeout Interval

Amount of time during which there is no activity, and after which a timeout occurs.

Transactions Processing Limit

Maximum number of LUWs (files) that are processed by the adapter in one request.

File Listener Information

Source File Type

Suffix of the listening adapter.

Source Directory (File)

Data source that is listened on.

File Extension

Extension of processed files.

Status

Current status of the adapter.

Processed Files

Number of processed LUWs (files).

Processed Bytes

Number of bytes processed by the file listener adapter.

Files in wait queue

Number of selected files waiting to be processed by the file listener adapter.

Processing time

Rounded number of seconds that the file listener adapter spent in processing files.

Wait for work time

Rounded number of seconds that the file listener adapter spent in waiting for files.

Average files per second

Average number of files per second that the file listener adapter processed.

Average bytes per second

Average number of bytes per second that the file listener adapter processed.

[Average file wait time](#)

Average time in seconds that the file listener adapter spent in waiting for files.

[Polling Interval](#)

Interval at which polling takes place.

[Timeout Interval](#)

Amount of time during which there is no activity, and after which a timeout occurs.

[Files Processing Limit](#)

Maximum number of LUWs (files) that are processed by the adapter in one request.

System

[CPU Time](#)

Total CPU time used by the process.

[Memory Usage \(KB\)](#)

Amount of memory, in kilobytes, used by the process.

[Disk Usage \(KB\)](#)

Disk space, in kilobytes, used by the process.

[Process ID](#)

Operating system specific process identification number that is associated with the agent.

Procedure: How to Display Statistics for an Individual Data Service Agent

1. Customize and display the Data Services Agents page, as described in [Displaying and Customizing a Data Service Agents Page](#) on page 661.
2. On the page, right-click the row for the individual Data Service Agent whose statistics you want to display. For example, right-click the row for the agent identified as Tscomid 11.
3. From the menu, click *Statistics*.

The Agent Statistics pane opens. On this pane, you will see the statistics described in [Data Service Agents Performance Parameters](#) on page 663, followed by system-extended statistics for the selected agent process.

Reference: System-Extended Statistics for an Individual Data Service Agent

The following table describes the system-extended statistics for an individual Data Service Agent.

Statistic	Description
PageFaultCount PeakWorkingSetSize (KB) WorkingSetSize (KB) QuotaPeakPagedPoolUsage QuotaPagedPoolUsage QuotaPeakNonPagedPoolUsage QuotaNonPagedPoolUsage PagefileUsage PeakPagefileUsage	Standard Windows memory management statistics for the monitored process.
Total Execution Time	Cumulative execution time of transactions performed by the agent.

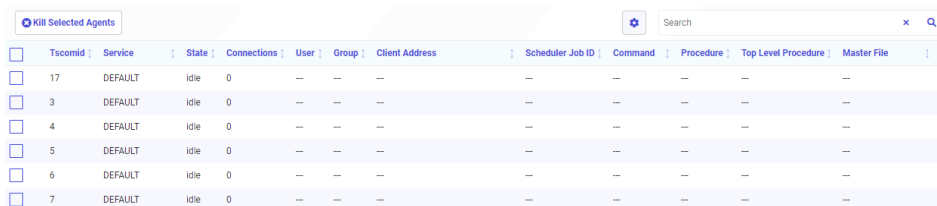
Displaying Status Statistics and Aggregated Statistics for Data Services

You can display and customize a status report for all Data Services, and display aggregated statistics for all Data Services or for a selected service.

To display and customize a Data Services Status Page:

1. Go to the *Workspace* page.
2. On the resources pane, right-click *Data Services*, and click *Services*.

The Data Services Status page is displayed, as shown in the following image.



<input type="checkbox"/>	Tscomid	Service	State	Connections	User	Group	Client Address	Scheduler Job ID	Command	Procedure	Top Level Procedure	Master File
<input type="checkbox"/>	17	DEFAULT	idle	0	--	--	--	--	--	--	--	--
<input type="checkbox"/>	3	DEFAULT	idle	0	--	--	--	--	--	--	--	--
<input type="checkbox"/>	4	DEFAULT	idle	0	--	--	--	--	--	--	--	--
<input type="checkbox"/>	5	DEFAULT	idle	0	--	--	--	--	--	--	--	--
<input type="checkbox"/>	6	DEFAULT	idle	0	--	--	--	--	--	--	--	--
<input type="checkbox"/>	7	DEFAULT	idle	0	--	--	--	--	--	--	--	--

The table in [Aggregated Statistics for Data Services](#) on page 671 describes all the statistics that you can display on the report.

3. To choose specific statistics that you want to display on the report, click *Choose Columns* from the *View* menu, and proceed according to the instructions in [Displaying and Customizing a Data Service Agents Page](#) on page 661.
4. Optionally, you can click the *Refresh* icon on the menu bar to refresh the statistics.

Procedure: How to Display Aggregated Statistics for All Data Services

1. Go to the *Workspace* page.
2. Click *Monitor*, then click *Data Service Statistics*, or right-click *Data Services* on the resources pane, and click *Statistics*.

The Data Service Statistics pane opens. The following image shows a sample Data Service Statistics pane.

Data Services Agents		Data Service Statistics	
General			
Service	All Services		
Performance			
Agents Now	7		
Agents Peak	5		
Connections Total	321		
Connections Completed	320		
Avg. Response Time	0:00:00.412		
Running Avg. Response Time	0:00:00.000		
Avg. DBMS Time	0:00:00.005		
Running Avg. DBMS Time	0:00:00.000		
Avg. Server Time	0:00:00.407		
Running Avg. Server Time	0:00:00.000		
Queued Now	---		
Queue Timeout	---		
Queued Peak	---		
Queue Time	---		
Errors			
Resource Fails	0		
Security Fails	0		
Focus Errors	7		
System			
CPU Time	0:00:27.138		
Memory Usage (KB)	180992		
Disk Usage (KB)	---		

The table in [Aggregated Statistics for Data Services](#) on page 671 describes the statistics that are provided.

Procedure: How to Display Aggregated Statistics for a Selected Data Service

1. Go to the *Workspace* page.
2. In the resources pane, expand *Data Services*.
3. Right-click the name of the service on which you want to report, for example, *WC_DEFAULT*, and click *Statistics*.

The aggregated statistics for the selected service are displayed on the right. The table in [Aggregated Statistics for Data Services](#) on page 671 describes the statistics that are provided.

Reference: Aggregated Statistics for Data Services

The following aggregated statistics are provided for Data Services:

General*Service*

If you are displaying aggregated statistics for all Data Services, this value is All Services.

If you are displaying aggregated statistics for a selected service, this value is the name of the service configured for Workspace Manager in the *edaserve.cfg* file.

Performance*Agents Now*

Number of agents that are currently running.

Agents Peak

Maximum number of agents that were running at any given time.

Connections Total

Number of attempted connections to servers since they were started, or from the time that statistics were reset.

This value is different from the total number of sessions, since a persistent session can connect multiple times for different requests.

Connections Completed

Number of successfully completed connections to servers since they were started, or from the time that statistics were reset.

Avg. Response Time

Average number of seconds for server response. This value does not include data transfer time. See [Calculating and Displaying Averages](#) on page 673, for more information.

Running Avg. Response Time

Response time, averaged over the time interval defined by the edaserve.cfg parameter stats_runningavg. The default value is 600 seconds.

Avg. DBMS Time

Average DBMS processing time in seconds.

Running Avg. DBMS Time

Average DBMS processing access time in seconds.

Avg. Server Time

Average response time minus DBMS processing time in seconds.

Running Avg. Server Time

Average response time minus DBMS processing access time in seconds.

Queued Now

Number of currently queued connections. Not applicable when queuing is off.

Currently running jobs are included in the connection total, but not in connections completed or in resource failures.

Queue Timeout

Number of connections that went into the queue and timed out (not available when queuing is off).

Queued Peak

Maximum number of connections that are queued at any given time (not available when queuing is off).

Queue Time

Average number of seconds spent in queue by all connections (not available when queuing is off).

Errors

Resource Fails

Number of connections rejected for lack of available agents. This value includes timed-out queued connections.

Security Fails

Number of connections that failed due to invalid credentials. For details about errors and failures, see the edaprint.log file, which is saved where the server is installed.

Focus Errors

Number of completed connections with WebFOCUS errors during request execution. For example, the adapter was not properly configured, or the data description or procedure was not found.

System

CPU Time

Total CPU time used by the server.

Memory Usage (KB)

Amount of memory, in kilobytes, used by the server.

Disk Usage (KB)

Total disk space, in kilobytes, used by the server.

Reference: Calculating and Displaying Averages

Variables that are averages (avg) are fractional numbers of seconds rounded to the nearest millisecond for display purposes, but actually computed in higher precision, depending on the operating system:

- Avg Rsp Time is a measure of waiting time - idle time + running (that is, agent processing) time for all connections, divided by the number of connections. It is actually the average time from the moment that the user clicks to send a request to connect or resume, until the moment that the answer appears in the browser, when the user suspends or disconnects the session.
- Agent processing time is the part of the connection duration spent only on running the request.
- Running time proportion of a connection is the percentage of its running time compared to its duration.

The accuracy of the three corresponding averages is limited only by the precision of the operating system. In the rare case in which a machine is faster than the precision of its time-measuring, accuracy side effects may occur:

- For an individual connection that has a duration shorter than the precision of the operating system, duration and running time cannot be measured and are both 0. The server considers it a 100% running time. A high occurrence of these instances produces an overestimated average proportion.

- ❑ If the duration is longer than the precision because of waiting time, but the running time is still less than the precision, then a zero is counted toward the average running time and in calculating the proportion. A 0% running time proportion is recorded. A high occurrence of these instances produces an underestimated proportion.

Working With the Connection Queue

The Connection Queue displays the statistics for the current list of connections, enabling administrators to monitor, and if necessary, cancel individual connections.

A connection refers to a physical connection between Client and Server. There are two types of connections:


- ❑ An active connection is one that is assigned to a session in a data access agent.
- ❑ A queued connection is one for which there are no available agents for the requested service, and the service is configured with a queue.

A queued connection waiting for an agent becomes active as soon as an agent is available. If the maximum time to wait in the queue is reached, the connection is automatically cancelled by the Workspace Manager.

***Procedure:* How to Display Statistics for All Connections**

1. Go to the *Workspace* page.

- On the resources pane, right-click *Data Services*, and click *Connection Queue*, as shown in the following image.



The Connection Queue report is displayed.

The table in [Connection Statistics](#) on page 675 describes the statistics on the report.

Each connection is represented as a separate row on the report.

Reference: Connection Statistics

The following statistics appear on the Connection Queue report.

Service

Name of the data service for the connection.

Type

Connection type, either connect or resume.

User

User ID that connected to the session.

Authentication

Type of authentication used to connect, either trusted, explicit, or IWA.

Requester ID

Unique identifier for the network connection created by a listener.

Time In

Time at which the connection was activated or queued. This value determines when a queued connection times out.

Client Address

Network address of the connecting client.

Foccache Ticket

Unique identifier assigned to the last request that used the foccache directory.

Procedure: How to Kill a Connection

To cancel a connection, click the row for the connection on the Connection Queue report, and from the menu, click *Kill Connection*.

- For an active connection, its session is forcefully disconnected from its agent.
- For a queued connection, it is simply cancelled and the client gets the same error that it would get if queuing was off and there were no available agents, or if the queue was full.

Performing Administrative Tasks for Data Services

You can perform certain administrative tasks as you monitor Data Service activity.

- Reset aggregated statistics for Data Services.
- Save the last FOCUS command.
- Kill one or more Data Service Agents.
- Kill agents by connected userid
- Start one or more new Data Service Agents.
- Display an error message for a Data Service Agent.

Procedure: How to Reset Aggregated Statistics for Data Services

You can clear aggregated Workspace Manager statistics for Data Services and start collecting them from the beginning.

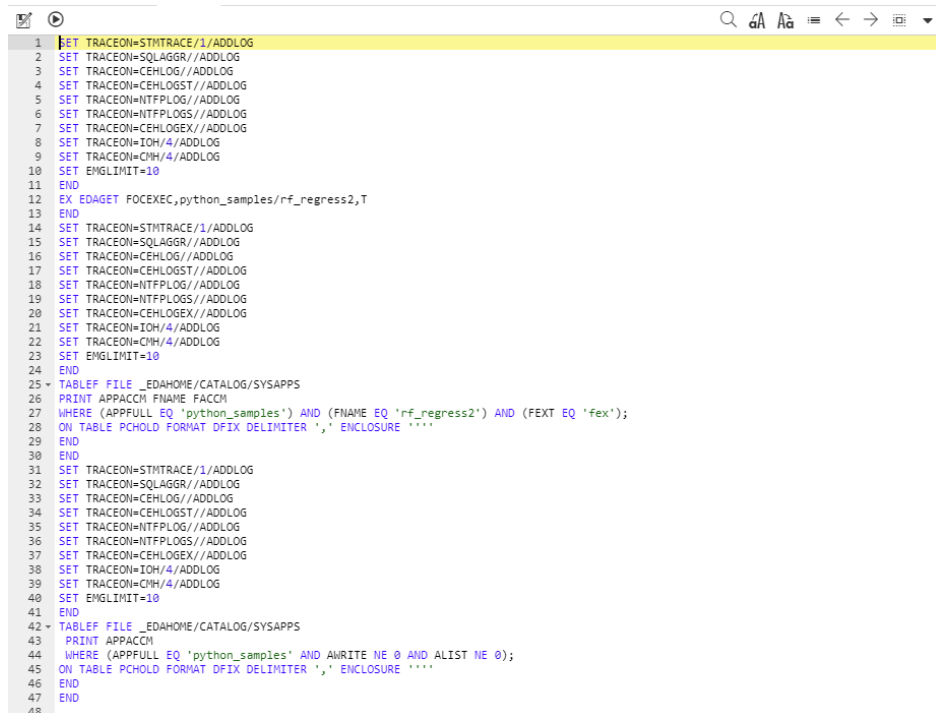
1. Go to the *Workspace* page.
2. In the resources pane, right-click *Data Services*.
3. From the menu, click *Reset Statistics*.

Procedure: How to Save the Last FOCUS Command

You can save the last FOCUS command executed by an agent as a new procedure and run it separately from the Reporting Server browser interface. This provides a great tool for debugging server issues.

1. Go to the *Workspace* page.
2. Right-click an agent on the Data Services Agent page and click *Save Last FOCUS Command*.

The text editor opens, as shown in the following image.



```

1  SET TRACEON=STHTRACE/1/ADDLOG
2  SET TRACEON=SQLAGGR//ADDLOG
3  SET TRACEON=CEHLOG//ADDLOG
4  SET TRACEON=CEHLOGST//ADDLOG
5  SET TRACEON=NTFPLOG//ADDLOG
6  SET TRACEON=NTFPLOGS//ADDLOG
7  SET TRACEON=CEHLOGEX//ADDLOG
8  SET TRACEON=IOH/4/ADDLOG
9  SET TRACEON=CH/4/ADDLOG
10 SET ENGLIMIT=10
11 END
12 EX EDAGET FOCXEC,python_samples/rf_regress2,T
13 END
14 SET TRACEON=STHTRACE/1/ADDLOG
15 SET TRACEON=SQLAGGR//ADDLOG
16 SET TRACEON=CEHLOG//ADDLOG
17 SET TRACEON=CEHLOGST//ADDLOG
18 SET TRACEON=NTFPLOG//ADDLOG
19 SET TRACEON=NTFPLOGS//ADDLOG
20 SET TRACEON=CEHLOGEX//ADDLOG
21 SET TRACEON=IOH/4/ADDLOG
22 SET TRACEON=CH/4/ADDLOG
23 SET ENGLIMIT=10
24 END
25 * TABLE FILE _EDAHOME/CATALOG/SYSAPPS
26 PRINT APPACM FNAME FACM
27 WHERE (APPFULL EQ 'python_samples') AND (FNAME EQ 'rf_regress2') AND (FEXT EQ 'fex');
28 ON TABLE PCHOLD FORMAT DFIX DELIMITER ',' ENCLOSURE *****
29 END
30 END
31 SET TRACEON=STHTRACE/1/ADDLOG
32 SET TRACEON=SQLAGGR//ADDLOG
33 SET TRACEON=CEHLOG//ADDLOG
34 SET TRACEON=CEHLOGST//ADDLOG
35 SET TRACEON=NTFPLOG//ADDLOG
36 SET TRACEON=NTFPLOGS//ADDLOG
37 SET TRACEON=CEHLOGEX//ADDLOG
38 SET TRACEON=IOH/4/ADDLOG
39 SET TRACEON=CH/4/ADDLOG
40 SET ENGLIMIT=10
41 END
42 * TABLE FILE _EDAHOME/CATALOG/SYSAPPS
43 PRINT APPACM
44 WHERE (APPFULL EQ 'python_samples' AND AWRITE NE 0 AND ALIST NE 0);
45 ON TABLE PCHOLD FORMAT DFIX DELIMITER ',' ENCLOSURE *****
46 END
47 END
48

```

3. Click the Save As button.

The Save As dialog box opens.

4. Select an application folder.
5. Enter a name in the *File Name* field.
6. Click *OK*.

Procedure: How to Kill One or More Data Service Agents

There are a number of ways to kill Data Service Agents:

- ❑ To stop one or a number of agents, select the check box next to each agent, and click *Kill Selected Agents*, as shown in the following image.

	Agent ID	Service	State	Connections	User	Group	Client Address	Scheduler Job ID	Command
<input checked="" type="checkbox"/>	11	WC_DEFAULT	in use	3			fe80::641f60b7:e7cc:3e56:53632	--	TABLE
<input type="checkbox"/>	12	WC_DEFAULT	idle	2	--	--	--	--	TABLEF
<input type="checkbox"/>	13	SCHED_DEFAULT	idle	0	--	--	--	--	--
<input type="checkbox"/>	14	SCHED_DEFAULT	idle	0	--	--	--	--	--
<input checked="" type="checkbox"/>	17	DEFAULT	idle	0	--	--	--	--	--

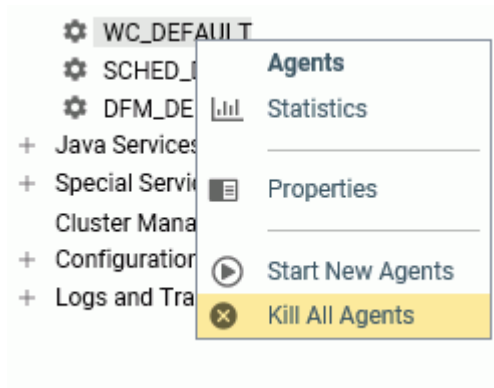
The running agent is terminated, invalidating any current connection to that agent. If a request is then issued from such a connection, an error message is returned. Stopping an agent is therefore an emergency administrative measure, as it disrupts the application. After an agent is terminated, the Data Services Agents pane refreshes automatically. The corresponding row of the terminated agent remains, and the State is listed as *stopping* until the row eventually disappears. You can also click *Refresh* on the tab header menu to update the pane and remove the row.

- ❑ To stop an individual agent, you can also right-click the row for the agent, and click *Kill This Agent* from the menu. When asked if you want to continue to kill the agent, click *OK*.
- ❑ To stop an individual agent without permitting the DBMS to perform any cleanup tasks, choose *Hard Kill This Agent* from the menu. A message informs you that a hard kill may leave resources inaccessible and should be used only as a last resort. Click *OK* to continue the hard kill.

Important: Use this option only when *Kill This Agent* fails.

- ❑ To stop all agents belonging to the same user ID, look for a row for the agent that is currently connected to the desired user ID. Right-click that row, and click *Kill Agents With The Same User ID*.

- ❑ To stop all agents associated with a selected service, right-click the service beneath the Data Services folder on the resources tree, and select *Kill All Agents* from the menu, as shown in the following image. When asked to confirm your request, click *OK*.



- ❑ For extreme situations in which you want to terminate all running agents, right-click *Data Services* on the resources tree. From the menu, click *Kill All Agents*. When asked to confirm your request, click *OK*.

Procedure: How to Kill Agents by Connected User ID

In the Reporting Server browser interface, look for a row for the agent in the Data Services Agents pane that is currently connected to the desired user ID, then right-click that row and click *Kill Agents With The Same User ID*, as shown in the following image.

<input type="checkbox"/>	4	DEFAULT	in use	PTH\svradmin	..1-50687	---
<input type="checkbox"/>	5	DEFAULT	idle	---	---	---
<input type="checkbox"/>	6	DEFAULT	idle	---	---	---
<input type="checkbox"/>	7	DEFAULT	idle	---	---	---
<input type="checkbox"/>	8	DEFAULT	idle	---	---	---
<input type="checkbox"/>	9	DEFAULT	idle	---	---	---
<input type="checkbox"/>	10	DEFAULT	idle	---	---	---
<input type="checkbox"/>	15	WC_DEFAULT	idle	---	---	---
<input type="checkbox"/>	13	SCHED_DEFAULT	idle	---	---	---
<input type="checkbox"/>	14	SCHED_DEFAULT	idle	---	---	---

To kill agents from the command line, the syntax is:

```
edastart -killuser <userid>
```


In addition, for the z/OS platform, a z/OS operator MODIFY command can also be used, as follows:

```
/F <server_jobname/started task>,-killuser <userid>
```

Procedure: How to Start One or More New Data Service Agents

Use this method to start additional agents, even when there are agents already running.

1. Go to the *Workspace* page.
2. On the resources pane, expand *Data Services*.
3. Right-click the name of the service for which you want to start agents.
4. From the menu, click *Start New Agents*.
5. On the Start New Data Service Agents pane, enter the number of new agents to start for the selected service, as shown in the following image. Then click *Start*.



On the Data Services Agents pane, you will see that the new agents are listed, with the State *starting*.

Procedure: How to Display an Error Message for a Data Service Agent

If a procedure generates an error, you can display the error message on the active agent.

1. Go to the *Workspace* page.
The Data Service Agents pane opens.
2. Right-click any column for an agent, and click *View Error* from the menu.
The latest error is displayed in a separate window.

Example: Displaying an Error Message for a Data Service Agent

The following error was generated on an agent when a procedure was run against a file that could not be found. Clicking the number of the agent that ran the procedure produced the following message:

```
(FOC205) THE DESCRIPTION CANNOT BE FOUND FOR FILE NAMED: XYZ  
BYPASSING TO END OF COMMAND
```

Tip: The error message is also stored in the Session Log, which you can access by clicking *Session Log* on the Reporting Server browser interface *User* menu. In the Session Log, you can click the message to see expanded information, including troubleshooting suggestions. For example:

```
(FOC205) THE DESCRIPTION CANNOT BE FOUND FOR FILE NAMED:
No data description can be found for the file requested. Check the
Master File Descriptions allocated or available, or the spelling of
the file name.
```

For more information on the Session Log, see [Showing the Session Log](#) on page 773.

Monitoring Java Services

From the Reporting Server browser interface, you can:

- Display a Performance Report for all Java Service Agents. The report displays statistics for the current list of agents monitored by the Workspace Manager. Multiple connections are supported, providing redundancy and load balancing capabilities. Also provided is the ability to refresh JSCOM3 instances after a specified number of connections, while ensuring that at least one connection is always available.
- Display statistics for an individual Java Service Agent.
- Display aggregated statistics for all Java Services.

As you monitor Java Service activity, you can also perform the following administrative tasks from the Reporting Server browser interface.

- Reset aggregated statistics for Java Services and start collecting them from the beginning.
- Stop or start one or more existing Java Service Agents.
- Kill one or more Java Service Agents.
- Start one or more new Java Service Agents.

Displaying and Customizing a Performance Report for Java Service Agents

This topic describes how Server Administrators can display a Performance Report for the agents for all Java Services or statistics for an individual Java Service Agent.

To display a Performance Report for All Java Service agents:

1. Go to the *Workspace* page.
2. In the resources pane, right-click the *Java Services* folder, and click *Agents*.

The Java Services Agents pane opens.

The following image shows a sample report.

Jscomid	Port	Status	Process ID	Connections	Active Connections	Last User	Last Executed Class	Free Memory (KB)
1	8123	active	12172	4	0	IBI\se02	ibi.callpgm.cpjmodel	497648

The table in [Performance Report for Java Service Agents](#) on page 682 describes the statistics displayed on the report.

3. At the upper-right of the resources pane, optionally click the Refresh icon to refresh the statistics, or click the *Auto-Refresh* icon and select the number of seconds between automatic refreshes.

Reference: Performance Report for Java Service Agents

The following statistics appear on the Performance Report for Java Service Agents. Each agent is represented as a separate row on the report.

Jscomid

Unique identification number assigned to each Java Service Agent. You can start and kill JSCOM3 connections. Each started connection is issued a unique number. If you kill a connection, its number is retired. A new connection is assigned an entirely new Jscomid.

Port

Port number for each Java Service Agent. Numbering starts with a base number and increases by one for each active agent.

Status

Current status of the instance. Possible values are starting, active, stopping, or stopped.

Process ID

Operating system specific process identification number associated with each Java Service Agent.

Connections

Total number of clients connected to JSCOM3, both active and inactive.

Active Connections

Active client connections to JSCOM3.

Last User

Name of the user who executed the class.

Last Executed Class

Last executed Java class. This includes explicitly called Java classes (such as CALLJAVA), and implicitly called classes (such as GRAPH FILE).

Free Memory (KB)

Amount of memory, in kilobytes, available to the Java Virtual Machine (JVM).

Procedure: How to Display Statistics for an Individual Java Service Agent

1. Display the Performance Report for all Java Service Agents, as described in [Displaying and Customizing a Performance Report for Java Service Agents](#) on page 681.
2. Right-click the row for the individual Java Service Agent whose statistics you want to display, and select *View Statistics*.

The Statistics pane opens, as shown in the following image. On this pane, you will see the statistics described in *Statistics for an Individual Java Service Agent* on page 685.

General	
Name	JSCOM3
Remark	---
State	active
Process ID	12172
Port Number	8123
System	
PageFaultCount	160275
PeakWorkingSetSize (KB)	383176
WorkingSetSize (KB)	263832
QuotaPeakPagedPoolUsage	252432
QuotaPagedPoolUsage	251984
QuotaPeakNonPagedPoolUsage	25912
QuotaNonPagedPoolUsage	24256
PagefileUsage	684539904
PeakPagefileUsage	724430848
Total Execution Time	0:00:10.484
Connections	
Jscomid	1
Number of Connections	4
Number of Active Connections	0
Last User	IBI\se02703
Last Executed Class	ibi.callpgrm.cpjmodel

3. At the upper-right of the resources pane, optionally click the Refresh icon to refresh the statistics, or click the *Auto-Refresh* icon and select the number of seconds between automatic refreshes.

Reference: Statistics for an Individual Java Service Agent

The following statistics are for an individual Java Service Agent.

General**Name**

Name of the Listener that is serving the Java Service.

Remark

Additional property of the Service or Listener, where available.

State

Current state of the connection. Possible values are starting, active, stopping, or stopped.

Process ID

Operating system specific process identification number associated with the Java Service.

Port Number

Port number for the Java Service. Numbering starts with a base number and increases by one.

System**PageFaultCount****PeakWorkingSetSize (KB)****WorkingSetSize (KB)****QuotaPeakPagedPoolUsage****QuotaPagedPoolUsage****QuotaPeakNonPagedPoolUsage****QuotaNonPagedPoolUsage****PagefileUsage****PeakPagefileUsage**

Standard Windows memory management statistics for the monitored process.

Total Execution Time

Cumulative execution time of transactions performed by the agent.

Connections

Jscomid

Unique identification number assigned to the Java Service Agent. You can start and kill JSCOM3 connections. Each started connection is issued a unique number. If you kill a connection, its number is retired. A new connection is assigned an entirely new Jscomid.

Number of Connections

Total number of clients connected to JSCOM3, both active and inactive.

Number of Active Connections

Total number of active client connections to JSCOM3.

Last User

Name of the user who executed the class.

Last Executed Class

Last executed Java class. This includes explicitly called Java classes (such as CALLJAVA), and implicitly called classes (such as GRAPH FILE).

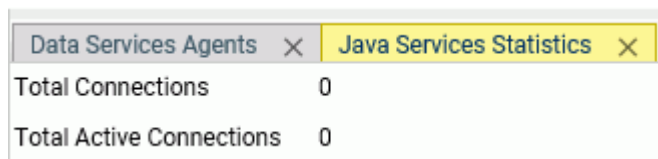
Displaying Aggregated Statistics for Java Services

You can display aggregated statistics for all Java Services.

To display aggregated statistics for all Java Services:

1. Go to the *Workspace* page.
2. On the resources pane, right-click *Java Services*, and click *Statistics*.

The Java Services Statistics pane opens on the right, as shown in the following image.



Data Services Agents	×	Java Services Statistics	×
Total Connections		0	
Total Active Connections		0	

Total Connections

Total number of clients connected to JSCOM3, both active and inactive.

Total Active Connections

Total number of active client connections to JSCOM3.

- On the menu bar on the navigation pane, optionally click the Refresh icon to refresh the statistics, or click the *Auto-Refresh* to turn on Auto-Refresh and enter the number of seconds between automatic refreshes.

Performing Administrative Tasks for Java Services

You can perform certain administrative tasks as you monitor Java Service activity.

- Reset aggregated statistics for Java Services.
- Stop or start one or more existing Java Service Agents.
- Stop or start all agents for a selected Java Service.
- Kill one or more Java Service Agents.
- Start one or more new Java Service Agents.

Procedure: How to Reset Aggregated Statistics for Java Services

You can clear aggregated Workspace Manager statistics for Java Services and start collecting them from the beginning.

- Go to the *Workspace* page.
- On the resources pane, right-click *Java Services*, and click *Reset Statistics*.

Procedure: How to Stop an Individual Java Service Agent

- Display the Performance Report for all Java Service Agents, as described in [Displaying and Customizing a Performance Report for Java Service Agents](#) on page 681.
- Right-click the row for the individual Java Service Agent that you want to stop, and click *Stop*, as shown in the following image.

Jscomid	Port	Status	Process ID	Connections	Active Connections
1	8123	active	8424	0	0

View Statistics

Stop

- When asked if you want to continue to stop this JSCOM3 instance, click *OK*.

On the Performance Report, you will see the status *stopping*. If you click *Refresh* on the tab header, you will see the status *stopped*.

Tip: You can restart the agent by right-clicking the applicable row on the Performance Report and clicking *Start* from the menu.

Procedure: How to Stop All Agents for a Selected Java Service

1. Go to the *Workspace* page.
2. On the resources pane, expand *Java Services*.
3. Right-click the service whose agents you want to stop, for example, *DEFAULT*, and click *Stop All*.
4. When asked if you want to continue to stop all Java Service Agents, click *OK*.

Tip: You can restart the agents by right-clicking the service on the resources tree, and clicking *Start All* from the menu.

Procedure: How to Kill One or More Java Service Agents

1. Go to the *Workspace* page.
2. On the resources pane, expand *Java Services*.
3. Right-click the name of the service for which you want to kill agents, and click *Kill Agents*.
4. On the Kill Java Service Agents pane, type the number of agents that you want to kill, and click *Kill*.

On the Performance Report for all Java Service Agents, you will see that the number of agents that you specified is removed from the list.

Procedure: How to Start One or More New Java Service Agents

Use this method to start one or more additional agents, even when there are agents already running.

1. Go to the *Workspace* page.
2. On the resources pane, expand *Java Services*.
3. Right-click the name of the service for which you want to start agents, and click *Start New Agents*.
4. On the Start New Java Service Agents pane, type the number of new agents to start for the selected service, and click *Start*.

On the Performance Report for all Java Service Agents, you will see that the new agents are listed, with the Status of *starting*, as shown in the following image.

Jscomid	Port	Status	Process ID	Connections	Active Connections
1	8123	active	8424	0	0
2		starting	15232	0	0
3		starting	2916	0	0

Monitoring Special Services and Listeners

From the Reporting Server browser interface, you can:

- Display a Performance Report for all Special Services and Listeners. The Performance Report displays statistics for the current list of Special Service processes and Listener processes monitored by the Workspace Manager.
- Display statistics for an individual Special Service or Listener.
- Display the Deferred List, which includes statistics and status for deferred requests.
- Display a Performance Report for all web sessions, or display statistics for an individual web session associated with an HTTP Listener.

As you monitor Special Service and Listener activity, you can also perform the following administrative tasks from the Reporting Server browser interface.

- Enable or quiesce certain Special Services or Listeners.
- Start or stop certain Special Services or Listeners.
- Kill all web sessions.
- View the log for a web session.

Procedure: How to Display a Performance Report for All Special Services and Listeners

1. Go to the *Workspace* page.
2. On the resources pane, right-click *Special Services and Listeners*, and click *Processes*.

The Performance Report for all Special Services and Listeners opens, as shown in the following image.

Name	Status	Process ID	Port Number / Directory	Number of Requests	CPU Time	Memory Usage (KB)
FDS	active	14188	8122	--	0:00:00.062	19200
SCHEDULER	active	13524	C:\ibi\srvt77\wfs\dfm_dir	0	0:00:00.078	15640
TCP/HTTP	active	10416	8120/8121	8/458	0:00:01.656	40108

The list in [Performance Report for All Special Services and Listeners](#) on page 690 describes the statistics displayed on the report.

3. On the resources tree header, optionally click the *Refresh* icon at the top of the resources tree. You can set a refresh interval by clicking *Auto-Refresh* at the top of the resources tree.

Reference: Performance Report for All Special Services and Listeners

The following statistics appear on the Performance Report for all Special Services and Listeners.

Name

Type of Special Service (for example, FDS or SCHEDULER), or type of Listener (for example, HTTP or TCP).

Status

Status of the Special Service or Listener. Possible values include active, not active, or stopped.

Process ID

Operating system specific process identification number associated with each Special Service or Listener running in a Workspace Manager address space.

Port Number/Directory

Port number or directory associated with the Special Service if applicable (for example, FDS port), or the port number associated with the Listener.

Number of Requests

Total number of requests processed by the Special Service or Listener.

CPU Time

Total CPU time used by the process, displayed as *hours:minutes:seconds.milliseconds* (for example, 0:00:05.625).


Memory Usage (KB)

Amount of memory, in kilobytes, used by the process.

Procedure: How to Display Statistics for an Individual Special Service or Listener

1. Display the Performance Report for all Special Services and Listeners, as described in [How to Display a Performance Report for All Special Services and Listeners](#) on page 689.

- Right-click the row for the individual Special Service or Listener whose statistics you want to display, and click *Statistics*, as shown in the following image.

Name	Status	Process ID	Port Number / Directory	Number of Requests	CPU Time	Memory Usage (KB)
FDS	active	14188	8122	--	0:00:00.062	19200
SCHEDULER	active	13524	C:\ibi\svr77\wfs\dfm_dir	0	0:00:00.078	15640
TCP/HTTP	active	 Statistics	0/8121	8/458	0:00:01.656	40108

The Statistics pane opens. On this pane, you will see the statistics described in [Statistics for an Individual Special Service](#) on page 691 or [Statistics for an Individual Listener](#) on page 694.

Tip: You can also display the Statistics pane for FDS, HTTP, TCP or SOAP by expanding *Special Services and Listeners* on the resources tree, right-clicking one of these Special Services or Listeners, and clicking *Statistics* from the menu.

Reference: Statistics for an Individual Special Service

The following statistics appear for an individual Special Service.

General

Name

Type of Special Service (for example, FDS, DFM, or SCHEDULER).

Remark

Additional property of the Special Service, where available.

State

Current state of the Special Service connection. Possible values are starting, active, stopping, or stopped.

Process ID

Operating system specific process identification number associated with the Special Service.

Port Number

Applies to the Special Service FDS.

Port number for the FDS. Numbering starts with a base number and increases by one.

Scanning

Applies to the Scheduler Special Services. Shows whether scanning is on or off.

System (Applies to the Special Services FDS and SCHEDULER)

PageFaultCount

PeakWorkingSetSize (KB)

WorkingSetSize (KB)

QuotaPeakPagedPoolUsage

QuotaPagedPoolUsage

QuotaPeakNonPagedPoolUsage

QuotaNonPagedPoolUsage

PagefileUsage

PeakPagefileUsage

Standard Windows memory management statistics for the monitored process.

Total Execution Time

Cumulative execution time of transactions performed by the agent.

Connections for FDS

Number of Connections

Total number of active connections.

Maximum Number of Connections

Hard-coded limit for the number of connections (512).

Number of FCB Blocks

Total number of opened File Control Blocks (FCB) by all users.

Maximum Number of FCB Blocks

Hard-coded limit for the number of opened FCBs at the sink side (4096).

Last User

User ID of the last connection to the FDS Service.

Last Executed HLI Commands

Last command executed by the FDS Service.

Connections for SCHEDULER**Status**

Current status of the Scheduler. For example, *scanning* means that the Scheduler is in the process of scanning the application path for scheduled jobs, and *ready* means that the Scheduler is ready to process a job.

Last Full Scan

Timestamp for the last full scan of the application path. A full scan is performed on the Scheduler *Start* and *Force Rescan* options, which are available on the Scheduler menu.

Last Incremental Scan

Timestamp for the last incremental scan. An incremental scan looks for files that have changed since the last scan. The last scan can be full or incremental, depending on which one happened last.

Number of Scheduled Events

Number of jobs currently scheduled to run.

Number of Scheduled Runs

Number of all runs since the startup of the Scheduler.

Number of CMASAP Runs

Number of all jobs submitted through CMASAP since the startup of the Scheduler.

DFM for SCHEDULER**DFM_DIR Available Disk Space (KB)**

Disk space, in kilobytes, in the directory `dfm_dir` that is available for the DFM Service write results of deferred requests.

Number of Requests Done Since Startup

Number of requests processed by the DFM Service since server startup.

Number of Responses Ready

Number of responses in the directory `dfm_dir` that are ready for a user to view (`get`).

Reference: Statistics for an Individual Listener

The following statistics appear for an individual Listener.

General

Name

Type of Listener, based on its network protocol (for example, HTTP or TCP).

Remark

Additional property of the Listener, where available. For example, an HTTP Listener can be Cluster Manager-enabled.

State

Current state of the Listener. Possible values include active, not active, and stopped.

Process ID

Operating system specific process identification number associated with the Listener.

Port Number

Port number for the Listener. Numbering starts with a base number and increases by one.

Total number of requests

Total number of requests processed by the Listener.

System

PageFaultCount

PeakWorkingSetSize (KB)

WorkingSetSize (KB)

QuotaPeakPagedPoolUsage

QuotaPagedPoolUsage

QuotaPeakNonPagedPoolUsage

QuotaNonPagedPoolUsage

PagefileUsage

PeakPagefileUsage

Standard Windows memory management statistics for the monitored process.

Total Execution Time

Cumulative execution time of transactions performed by the agent.

Setting a Web Session Timeout Warning

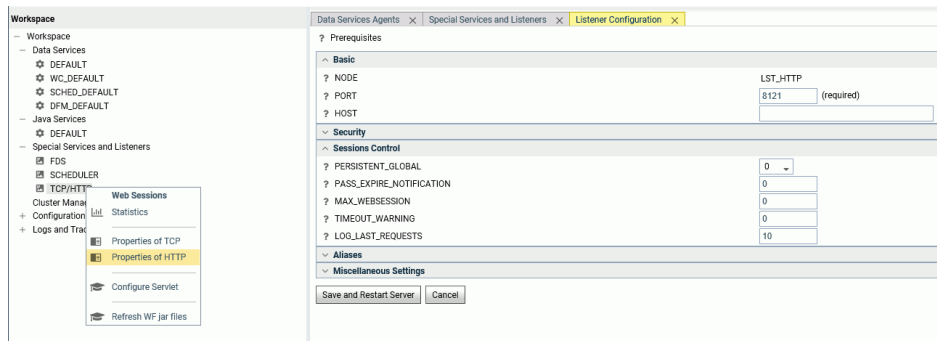
The server can post a warning about web session expiration before the session actually expires.

To set the number of minutes prior to session expiration that the warning displays:

1. On the Workspace resources tree, expand *Special Services and Listeners*.
2. Right-click TCP/HTTP and click *Properties of HTTP* from the shortcut menu.

The Listener Configuration page opens.

3. Expand the *Sessions Control* group, as shown in the following image.



4. Enter a number of minutes, n , in the TIMEOUT_WARNING text box.

By default, the value is zero (0), meaning that no warning will be displayed.

5. Click *Save and Restart Server*.

The `TIMEOUT_WARNING= n` attribute is added to the `odin.cfg` file.

Responding to the warning message resets the session expiration timer.

Deferred Lists

The Deferred List displays the current status of all deferred requests. You can stop the queuing or execution of deferred requests, and remove requests from the queue.

To display the Deferred List:

1. Go to the *Workspace* page.
2. On the resources pane, expand *Special Services and Listeners*.
3. Right-click *SCHEDULER*, and click *Deferred List*.

The Deferred List is displayed, as shown in the following image.

<input type="checkbox"/>	Defer ID	User ID	Status	App	Procedure	Report	Submitted Time	Scheduled Time	Days Left
<input type="checkbox"/>	20180328173953_3c1da145_LOOPBACK	PTH\sradmin	Ready	doc77x	dmdebugger3		2018 March 28 17:39:54	2018 March 28 17:39:54	29
<input type="checkbox"/>	20180328174018_717d836f_LOOPBACK	PTH\sradmin	Unknown	doc77x	crwfproduct		2018 March 28 17:40:19	2018 March 28 17:40:19	29

Each row displays statistics for a procedure that has been submitted for deferred execution. For details about these statistics, see [Deferred List Statistics](#) on page 696.

- All deferred requests are displayed by default. However, you can select the check box next to one or more deferred requests, then click *Delete Selected*, to remove those requests from the report.

Optionally, you can:

- Click the check box in the heading row to select the check boxes for all deferred requests.
 - Click *Refresh Selected* to refresh the status of the deferred request or requests whose check boxes are selected.
- The right-click menu for a row has the following options.
 - View, to view the output.
 - Download, to download the output to the Downloads directory on your PC.
 - Delete, to delete this row.

Reference: Deferred List Statistics

The following statistics appear on the Deferred List.

[Defer ID](#)

Unique identifier that the server assigns to each deferred request submitted.

[User ID](#)

ID of the user who requested deferred execution of the named procedure.

[Status](#)

Can be one of the following: Ready, Queuing, Executing, Agent crashed, Connect failed, Exceed limit, or Unknown.

[App](#)

Name of the application that contains the procedure being run in deferred mode.

Procedure

Name of the procedure being run in deferred mode.

Report

Name of the report being run in deferred mode.

Submitted Time

Time at which the request for deferred execution was submitted.

Scheduled Time

Time at which the deferred request is scheduled to run.

Days Left

Number of days that the generated report will be kept on the server.

Procedure: How to Retrieve the Answer Set for a Deferred Request

1. On the Deferred List, right-click a row for a deferred request whose status is Ready.
2. From the menu, click *View* to retrieve the answer set for the deferred request in that row.
The output is displayed.

Procedure: How to Delete a Deferred Request From the Queue

Click anywhere in a row for a deferred request, and from the menu, click *Delete*.

Tip: To delete multiple rows, click the check box next to the rows that you wish to delete. Then click the *Delete Selected* button.

Web Sessions

You can display a Performance Report for all web sessions and statistics for a single web session.

To display a Performance Report for all web sessions:

1. Go to the *Workspace* page.
2. On the resources pane, expand *Special Services and Listeners*, right-click *TCP/HTTP*, and click *Web Sessions*.

The Performance Report for all web sessions opens.

The list in [Performance Report for All Web Sessions](#) on page 698 describes the statistics displayed on the report.

3. Click *Refresh* at the top of the navigation pane to refresh the statistics, or click *Auto-Refresh* above the resources tree and select the number of seconds between automatic refreshes.

Reference: Performance Report for All Web Sessions

The following statistics appear on a Performance Report for all web sessions.

Web Session Number

Unique number assigned to a web session at logon to the server until the end of the session. The session ends by logging off the server or by closing the browser.

Web Session ID

A unique number that the server assigns to a specific user for the duration of the user session. Every time you connect to the Reporting Server browser interface, a new session number is assigned.

User

ID of the connected user.

Remote Address

IP address or machine on which the web browser running the Reporting Server browser interface resides.

First Access Time

Time at which the web session originated.

Last Access Time

Last time at which the user was active in this web session.

Authentication

Mechanism that performs user authentication at the time the web session is created.

Console Message Count

Number of messages that a user received in this session.

Procedure: How to Display Statistics for an Individual Web Session

1. Display the Performance Report for all web sessions, as described in [Web Sessions](#) on page 697.
2. Right-click the row for the individual web session whose statistics you want to display, and select *View Statistics*.

The Web Session Statistics pane opens.

The list in *Statistics for an Individual Web Session* on page 699 describes the statistics displayed on the pane.

Reference: Statistics for an Individual Web Session

The following statistics appear for an individual web session.

Web Session Number

Unique number assigned to a web session at logon to the server until the end of the session. The session ends by logging off the server or by closing the browser.

Session ID

A unique number that the server assigns to a specific user for the duration of the user session. Every time that you connect to the Reporting Server browser interface, a new session number is assigned.

User ID

ID of the connected user.

Role(s)

Roles of the connected user

Number of HTTP Requests

Number of HTTP requests processed during the current web session.

Number of TCP Requests

Number of TCP requests processed during the current web session.

First Access Time

Time at which the web session originated.

Last Access Time

Last time at which the user was active in this web session.

Remote Address

IP address or machine on which the web browser running the Reporting Server browser interface resides.

User-Agent

Standard HTTP client that is used for connection to the HTTP Listener.

Login Type

Type of user logon, for example, explicit or cookie.

Last 10 HTTP Request

List of URLs, Times, and Post Data for the last 10 HTTP requests.

Performing Administrative Tasks for Special Services and Listeners

You can perform several administrative tasks for Special Services and Listeners.

- Enable or quiesce certain Special Services or Listeners.
- Start or stop certain Special Services or Listeners.
- Kill all web sessions.
- View the log for the web session.

Procedure: How to Stop a Special Service or Listener

Stop is only available for certain special services.

1. Expand *Special Services and Listeners* on the resources tree.
2. Right-click a Special Service or Listener.
3. Click *Stop* to stop the service

The Performance Report or statistics pane reflects the stopped status. A *Start* option is added to the right-click menu for the service.

Procedure: How to Kill All Web Sessions

1. Display the Performance Report for all web sessions, as described in [Web Sessions](#) on page 697.
2. On the menu bar, click *Kill All Sessions*, as shown in the following image.

The screenshot shows a window titled 'Web Sessions' with a menu bar containing 'Kill All Sessions'. Below the menu bar is a table with the following data:

Session Number	Session ID	User	Group(s)	Remote Address	First Access Time
5	5abce26f:1-5	--	--	fe80::641f:60b7:e7cc:3e56%2:55427	2018 March 29 08:56:26
7	5abce26f:1-7	--	--	172.30.234.75:55921	2018 March 29 09:35:37

Procedure: How to View the Log for a Web Session

1. Display the Performance Report for all web sessions, as described in [Web Sessions](#) on page 697.
2. Right-click the row for the web session whose log you want to view.
3. From the menu, click *View Session Log*.

For more information, see [Showing the Session Log](#) on page 773.

Editing Configuration Files

Manually editing configuration files requires extensive knowledge of the inner workings of the server. All files except VERSION.CFG are available for editing in the Reporting Server browser interface text editor. VERSION.CFG is created by the installation/configuration process and is only available for viewing from that page.

To edit a configuration file:

1. Go to the *Workspace* page.
2. On the resources pane, expand *Configuration Files*. You can open the following files for editing:

- Workspace.** Opens edaserve.cfg.
- Administration.** Opens admin.cfg.
- Shared Resources.** Opens shared.cfg.
- Communication.** Opens odin.cfg.
- Version.** Opens version.cfg for viewing only.
- Server Profile.** Opens edasprof.prf.
- Trace Profile.** Opens ibitrace.fex.
- FDS Profile.** Opens suprof.prf.
- User/Group Profiles Folder.** Has an option for each defined user and group.

Note that both Server Administrators and Application Administrators can edit the *user.prf* file. However, Server Administrators can edit all *user.prf*s, while Application Administrators can only edit their own profiles.

- Miscellaneous folder.** Has the following options:
 - Server Startup Script.** Opens edastart.bat
 - LOCALE.** Opens nlscfg.err
 - Font.** Opens fontuser.xml

- ❑ **Environment.** Opens `edaenv.cfg` (When the server starts, environment variables are set according to entries in the `edaenv.cfg` file. If the environment conflicts with an entry in `edaenv.cfg`, the entry in `edaenv.cfg` takes precedence.)

Configuring Java Services

You can configure Java Services to add the JSCOM3 listener, which is required to support and monitor Java-based adapters and functionality, such as GRAPH. You can fine tune server performance.

The server has an automatic Java lookup feature on Windows to enable the Java Listener to start without any direct server configuration steps.

At server startup time, if no explicit `JAVA_HOME` or `JDK_HOME` variables are set either at the system environment level or in the server environment configuration file (`EDACONF\bin\edaenv.cfg`), the Windows registry will be searched for an applicable Java JDK or JRE release (currently 1.8 or higher), and the appropriate variable will be set.

This feature assumes that a Java installation has been done using the Oracle Windows JRE or JDK installation tool (which sets the Windows registry) prior to running the server (but the server itself requires no configuration).

To Configure Java Services for a JSCOM3 Listener:

1. Go to the *Workspace* page.
2. In the resources pane, expand *Java Services*.
3. Right-click *DEFAULT* and click *Properties*.

The Java Services Configuration pane opens, as shown in the following image.

^ Basic	
? NODE	JSS
? PORT	<input type="text" value="8123"/> (required)
? HOST	<input type="text"/>
? REFRESH	<input type="text"/>
? NUMBER_READY	<input type="text"/>
v Advanced	
v JVM Settings	
v Class path	
v Version and path	

4. Enter values for each of the configuration sections.
5. When finished, click *Save and Restart Java Services*.

Reference: **Java Services Configuration Parameters**

You can enter values on the Java Services Configuration pane for Basic and Advanced settings, JVM Settings, and the Classpath. For information on JVM Settings, see [How to Fine-Tune JVM Settings From the Reporting Server browser interface](#) on page 706.

Basic

NODE

Defines the logical name for the node block.

PORT

Defines the TCP port number that the listener is listening on.

HOST

Defines the IP address that the listener is listening on.

REFRESH

Defines the minimum number of new connections which can be accepted during the life of each Java server process. A Java server which has reached the limit will be terminated, and a new Java server will be started. To use this feature, you must configure more than one Java server (using NUMBER_READY) to prevent service gaps. It is ignored if only one Java server is configured.

NUMBER_READY

Defines the number of Java servers that the Workspace Manager will create at startup. The first server will listen on the PORT defined in the node. Each additional server will listener on an incremental port, for example, PORT + 1, PORT + 2, etc.

Advanced

GPOOLSIZE

Defines the number of Graph Service Agents that each Java server will create at startup.

FLEX_SUPPORT

Enables JAVA environment to support Flex. It should be set to "0" ONLY if there is a conflict between the CLASSPATH objects.

Class path

CLASSPATH

The environment variable that tells the JVM or Java-based applications where to look for user-defined classes and packages.

IBI_CLASSPATH

Defines the additional Java Class directories or full path jar names which will be available for Java Services.

Fine-Tuning the TIBCO WebFOCUS Reporting Server

You can fine-tune the server to optimize performance.

A Java Services folder on the Workspace Configuration resources tree enables you to create and manage multiple client connections to JSCOM3, thereby supporting redundancy and load balancing, as well as the ability to refresh JSCOM3 instances while ensuring that at least one connection is always available.

Java exceptions are displayed in the edaprint.log file, and JVM version and path information is displayed in the edaprint.log file and on the Java Services Configuration pane.

Using JVM Tuning Options to Improve Memory Usage

You can obtain better memory usage performance in production servers by properly tuning the JVM heap and stack sizes using three options on the JVM Settings tab of the Java Services Configuration pane. These options set initial Java heap size, maximum Java heap size, and Java thread stack size. The options update the `odin.cfg` file. For related information, see [Editing Configuration Files](#) on page 701.

Before setting these parameters, it is useful to understand the following relationships.

Heap Size

JSCOM3 is the server interface to the JVM. When you start JSCOM3, you can specify memory allocation for the JVM using the `-X` options shown in the following table.

JVM option passed to JSCOM3	Description
<code>-Xms</code>	Determines the initial (minimum) Java heap size.
<code>-Xmx</code>	Determines the maximum Java heap size

Heap size does not determine the amount of memory your process uses. As you monitor your Java process, you may notice that the amount of memory used exceeds the amount you have specified for `-Xmx`. This occurs because `-Xmx` limits the Java heap size and may allocate memory for other things, including a stack for each thread. Therefore, the total memory consumption of the JVM may exceed the value of `-Xmx`.

Stack Size

Each thread in the JVM is assigned a stack. The stack size limits the number of threads you can use. If the stack size is too large, you risk running out of memory since each thread is allocated more memory than it needs.

The following chart describes the JVM Option passed to JSCOM3.

JVM option passed to JSCOM3	Description
<code>-Xss</code>	Determines the stack size for each thread. You can make this value as low as 1K: <code>-Xss1K</code> . However, if the stack space is too small, you will eventually see an exception class <code>java.lang.StackOverflowError</code> .

Controlling Heap and Stack Size

JSCOM3 uses default values for the minimum (initial) and maximum heap sizes and for the stack size of each thread. To improve performance, you can override the JSCOM3 default settings using the initial Java heap size, maximum Java heap size, and Java thread stack size options. For details, see [How to Fine-Tune JVM Settings From the Reporting Server browser interface](#) on page 706.

You can also choose the *Force Finalization* option on the Java Services menu, which assists in memory management by releasing memory that is not currently in use by the Java Virtual Machine (JVM).

Procedure: How to Fine-Tune JVM Settings From the Reporting Server browser interface

1. Go to the *Workspace* page.
2. In the resources pane, open the *Java Services* folder, right-click a service, and select *Properties*. The Java Services Configuration pane opens.
3. Click the JVM Settings tab and enter values for the following parameters:

Initial Java Heap Size

Defines the initial Java heap size in megabytes (MB), which reflects the minimum amount of memory that is already being used.

Supply a positive integer number.

Maximum Java Heap Size

Defines the maximum Java heap size in megabytes (MB).

Supply a positive integer number.

Note: With server-side Java applications like JSCOM3, it is good practice to set the initial (minimum) and maximum heap sizes to the same value.

Java Thread Stack Size

Defines the Java thread stack size in kilobytes (KB).

Supply a positive integer number.

JVM_OPTIONS

Defines the additional JVM options. This is the parameters that can affect the performance characteristics of the JVM.

JVM_PROPERTIES

Defines the additional JVM properties. This is a platform-independent generalization of the system environment variables.

4. Click *Save and Restart Java Services* to update the `odin.cfg` file and restart the server with the new settings.

Cluster Manager

Cluster Manager enables you to balance workload across a number of servers. It has polling and monitoring capabilities to identify which servers are up-and-running, average query response time, and the number of concurrent tasks running. Cluster Manager can also identify which remote servers are down, and appropriately re-route requests, as well as route a user connection to the most efficient remote server to enhance overall performance.

While a server configured as a Cluster Manager Server can execute report requests and act as a Reporting Server, the Cluster Manager Server is generally used to distribute tasks to Reporting Servers within its configuration. For clarity, a server configured for Cluster Manager use is referred to as a Cluster Manager Server (or a Cluster Manager Cluster), while non-Cluster Manager Servers are simply Reporting Servers (or servers).

There are three forms of Cluster Manager Server configuration:

- Single (Classic) Cluster Manager Server.
- Multi-Cluster Manager Server with Dynamic Load Balancing. This configuration is designed for scalable cloud environments, but can also be used for on-premise environments. It requires external load balancing software (such as NGNIX) and ETCD (for inter-communication of server state and status) to be installed, running, and accessible.
- Multi-Cluster Manager Server On-Premise. This configuration is designed for on-premise environments with clustering, based on business needs. It requires no external software.

An additional feature is *Auto Cluster Manager Self Registration* of Reporting Servers to a respective Cluster Manager Server. Self registration is configured by adding the `clm = url` configuration parameter to a Reporting Server configuration, so the server automatically registers itself to the Cluster Manager Server when the Reporting Server is started. This allows fluid scaling of any number of Reporting Servers, without the need to explicitly configure the participants on the Cluster Manager Server.

If you are using Multi-Cluster Manager with Dynamic Load Balancing mode, the first step is to install and start your load balancing software and ETCD. They do not need to be on the machine where the server is running, nor do the balancer and ETCD need to be on the same machine, but you will need to know the machine node names and the ports on which they (the balancer and ETCD) are listening, and have a read/write access ID and password for ETCD. The next step (regardless of the mode you want to use) is to enable Cluster Manager.

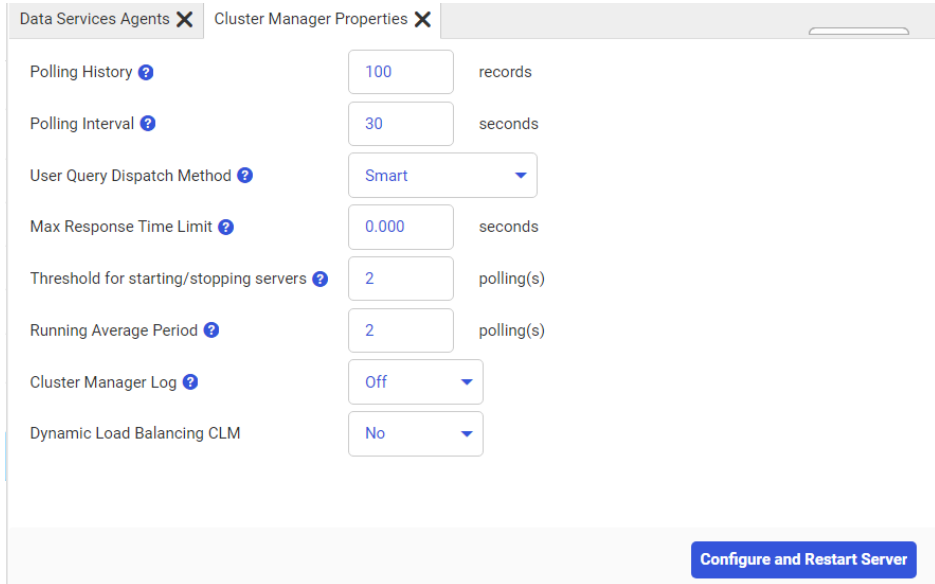
You can enable Cluster Manager, add remote servers, and configure clusters from the Workspace page.

Note: Cluster Manager connections to a remote server do not use the queuing capabilities that may be configured on remote servers that are members of a cluster.

To enable Cluster Manager:

1. Go to the *Workspace* page.
2. Right-click *Cluster Manager* and click *Enable*.

The Enable Cluster Manager pane opens, as shown in the following image.




The screenshot shows the 'Cluster Manager Properties' configuration pane. It contains several settings:

- Polling History: 100 records
- Polling Interval: 30 seconds
- User Query Dispatch Method: Smart
- Max Response Time Limit: 0.000 seconds
- Threshold for starting/stopping servers: 2 polling(s)
- Running Average Period: 2 polling(s)
- Cluster Manager Log: Off
- Dynamic Load Balancing CLM: No

A 'Configure and Restart Server' button is located at the bottom right of the pane.

Either accept the default values, or edit them. For information about Cluster Manager properties, see [Cluster Manager Properties](#) on page 719.

3. If you want to configure a Dynamic Load Balancing Cluster Manager, select Yes from the drop-down list, and the following additional input boxes display.



The following input boxes are displayed for ETCD configuration:

- ETCD URL
- ETCD User
- ETCD Password
- ETCD Key

Enter values for the following parameters.

ETCD URL

Enter node and port in the format `http://node:port`, or `https://node:port`, depending on whether your ETCD configuration is configured for SSL.

ETCD Name

Enter an ETCD user ID that has Read/Write access.

ETCD Password

Enter the password associated with the user ID entered in the ETCD Name entry field.

ETCD Key

Enter an arbitrary string (for example, Accounting). This value acts as a key, and allows the ETCD Server to be reused for multiple groups of Multi-Cluster Manager Servers, while eliminating cross-cluster noise between the unrelated Cluster Manager Servers. If you are only using a single group of Cluster Manager Servers in your ETCD environment, you can leave this field blank. However, supplying a value provides added identification when examining ETCD records, so it is highly recommended to supply a value.

4. Click *Configure and Restart Server*.
5. Click *OK* in the confirmation message.

The Workspace will restart.

6. For Multi-Cluster Manager Server configurations within a group, repeat the Cluster Manager enablement step for as many Cluster Manager Servers as you want to participate in the Cluster Manager group.

After each Cluster Manager Server configure and restart, you have the option to:

- Sign in to each of the Reporting Servers that you want to participate in the Cluster Manager Cluster, and configure *Auto Cluster Manager Self Registration*.
- Sign in to the Cluster Manager Server, and manually configure each of the Cluster Manager Servers for TCP connections to each of the Reporting Servers that you want to participate in the Cluster Manager Cluster. This is done on the Cluster Manager Server by navigating to *Workspace/Cluster Manager/New Remote Server*, and adding TCP connections for each of the Reporting Servers that you want to participate in the Cluster Manager Cluster.

Multi-Cluster Manager with Dynamic Load Balancing mode requires the client (WebFOCUS Mid-tier) side to have its `REMOTE_CLM_HOST` and `REMOTE_CLM_HOST` parameters point to the load balancer that is in use. Also, the balancer must point to the Cluster Manager Servers that are participating in the configuration.

Adding Auto Cluster Manager Self Registration to a Remote TIBCO WebFOCUS Reporting Server

You can configure Auto Cluster Manager Self Registration in one of the following ways:

- Enter the URL to the Cluster Manager Server in the Reporting Server browser interface Cluster Manager parameter entry field.
- Edit the reporting server edaserve.cfg file to add a parameter line pointing to a Cluster Manager Server URL.

The added parameter specifies the URL of the Cluster Manager Server so it can request registration to the cluster. For Multi-Cluster Manager Dynamic mode, the steps are the same. However, the parameter value is a comma-delimited list of Cluster Manager Server URLs instead of a single URL.

With either method, saving the configuration will force the server to recycle, so be prepared to implement this configuration change during off hours.

Each method is described in the following sections.

Procedure: How to Add Auto Cluster Manager Self Registration Using the Workspace Miscellaneous Settings Page

1. Sign in to the remote Reporting Server.
2. Navigate to the Workspace page, click *Settings*, then *Workspace SET*, then *Miscellaneous Settings*.
3. Scroll down to the Cluster Manager parameter and add the following in the text entry field.

- For a single Cluster Manager registration:

`http://host:port`

- For Multiple Cluster Manager registrations:

`http://host:port, http://host:port, ...`

where:

`host`

Is the host name where the Cluster Manager Server is running.

`port`

Is the HTTP port number on which the Cluster Manager Server is listening.

4. Click Save.

The server will recycle. Upon recycle, the server will register itself with the Cluster Manager and participate in the cluster.

Repeat these steps for each Reporting Server that you want to participate in the cluster.

The Cluster Manager panel (of the Cluster Manager Server) will reflect the current nodes participating in the cluster. For more information, see [Cluster Manager Global Monitoring Console](#) on page 721.

***Procedure:* How to Add Explicit Remote TIBCO WebFOCUS Reporting Servers**

Explicitly added Remote Servers can be used regardless of whether a Cluster Manager Server has Reporting Servers participate using the Auto Cluster Manager Self-Registration feature. However, a given Report Server should not use both methods to participate.

To add remote servers, Cluster Manager must be enabled.

1. Go to the *Workspace* page.
2. Right-click *Cluster Manager*, point to *New Remote Server*, and click *TCP*.

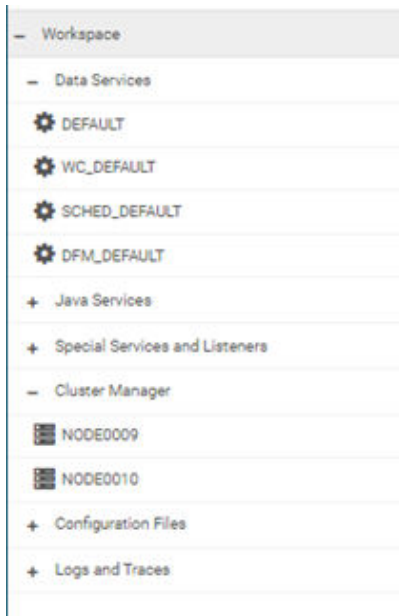
The Remote Server Configuration pane opens, as shown in the following image.

3. Enter the host name or IP address for the remote Reporting Server in the *HOST* field.
4. Enter the remote Reporting Server port in the *PORT* and *HTTP_PORT* fields.
5. Select a security mode from the *SECURITY* drop-down menu. The options are Trusted, Explicit, or IWA. Explicit will require you to enter a user and password, and a domain for servers running on Windows.
6. Optionally, enter a description for the Reporting Server in the *DESCRIPTION* field.
7. If you are adding a remote Reporting Server to start remotely on Windows, enter the Windows User, Windows Password, and Windows Service.

For more information, see [Remote TIBCO WebFOCUS Reporting Server Configuration Parameters](#) on page 713.

8. Optionally, expand the Advanced section to set additional parameters. For information, see [Remote TIBCO WebFOCUS Reporting Server Configuration Parameters](#) on page 713.
9. Click Save.
10. Repeat the procedure to add all the Reporting Servers that will be part of the cluster.

The Reporting Servers will be added in the Cluster Manager folder, as shown in the following image.



Reference: [Remote TIBCO WebFOCUS Reporting Server Configuration Parameters](#)

The Remote Reporting Server Configuration pane has Basic and Advanced sections.

Basic

NODE

Assigns a name to the remote Reporting Server. The node name is a character string of up to 130 characters.

HOST

Indicates the host name or IP address for the remote Reporting Server.

PORT

Defines the TCP port number for the Reporting Server.

HTTP_PORT

Defines the HTTP port number for the Reporting Server.

SECURITY

Defines the security mode. The options are Trusted, Explicit, or IWA. Explicit will require you to enter a user and password, and a domain for servers running on Windows.

Windows User, Windows Password, Windows Service

Required for remote start of a remote Reporting Server on Windows. It will start the remote Reporting Server using the indicated service with the user ID and password provided in these parameters. This feature is available for Cluster manager on Windows, to start Windows remote Reporting Servers.

Windows User

Is the Windows account user ID for a Reporting Server on Windows.

Windows Password

Is the Windows account password for a Reporting Server on Windows.

Windows Service

Is the Windows service name for a Reporting Server on Windows.

DESCRIPTION

Defines a description for the Reporting Server.

Advanced

SERVICE NAME

Indicates the name of the data service used by the remote Reporting Server.

COMPRESSION

Activates data compression for data transfer between client and Reporting Server.

ENCRYPTION

Defines the encryption method used in TCP data transfer between client and Reporting Server. The options are DES, Advanced, or IBCRYPT.

CONNECT_LIMIT

Defines the maximum time that the client will wait for a TCP connection response from the Reporting Server. The default value is -1, which creates an indefinite wait.

MAXWAIT

Defines the time that the client will wait for a response from the Reporting Server in seconds. The default value is -1, which creates an indefinite timeout

REMOTE_EDASTART

Defines the full name of the edastart command for the remote Reporting Server.

REMOTE_USER_PROFILE

Determined if all commands in the personal \$HOME/.profile will be executed before executing the edastart command for the remote Reporting Server. The default value is y.

TYPE

Defines the type of node. Reserved node will be brought up and down by Cluster Manager to optimize performance

Add to Cluster Node

Enables you to add the Reporting Server to a cluster.

Procedure: How to Configure Additional Cluster Manager Clusters

Enabling a Cluster Manager Server automatically configures a default cluster on the server. To configure additional remote Cluster Manager Clusters for the server:

1. Go to the *Workspace* page on the Cluster Manager Server.
2. Right-click *Cluster Manager* in the resources pane, point to *New Remote Server*, and click *CLUSTER*.

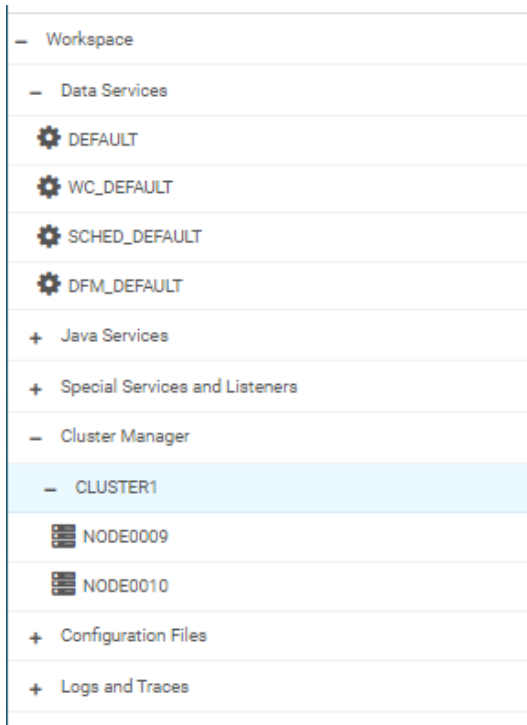
The Remote Cluster Configuration pane opens, as shown in the following image.

3. Enter an easily identifiable name for the cluster in the *NODE* field, for example, ACCOUNTING. The node name is a character string of up to 130 characters.
4. Click *SCHEDULER* to manage the Scheduler.

Choose this option if your Cluster is designated for the Deferred jobs and Scheduled jobs failover feature.

5. If you are using the Scheduler, select Explicit or Trusted security.
This security setting will be used to manage the Scheduler on remote Reporting Servers.
6. Select the check boxes for the remote Reporting Servers that you want to participate in the cluster.
7. Optionally, enter a description for the cluster in the *DESCRIPTION* field. If you leave it blank, the node name will be used for display purposes.
8. Click Save.
9. Optionally, expand the Advanced section to set additional parameters. For information, see [Remote Cluster Configuration Parameters](#) on page 716.

The cluster will be added in the Cluster Manager folder, as shown in the following image.



Reference: Remote Cluster Configuration Parameters

The Remote Cluster Configuration pane has Basic and Advanced sections.

Basic

NODE

Assigns a name to the cluster.

SCHEDULER

Click to manage the Scheduler.

SECURITY

Applies to the Scheduler. Select Explicit (the user must enter credentials) or Trusted (credentials used are those of the operating system user impersonated by the server data access agent)

ALTERNATE

Lists the servers that are available for the cluster.

DESCRIPTION

Defines a description for the cluster.

Advanced**CLM_MAXWAIT**

Defines the time that the client will wait for a response from the Cluster Manager in seconds. The default value is -1, which creates an indefinite wait.

DISPATCH_METHOD

Defines how Cluster Manager dispatches user queries to remote servers. The options are:

- Adhere to the settings in Cluster Manager.** The cluster will use the option set in the Cluster Manager settings. This is the default value.
- RANDOM.** The query will be dispatched randomly to a running server in the cluster. Every server will have the same weight. The query will be queued if a queue is available, and there is no available agent in the selected server.
- ROUND-ROBIN.** The query will be dispatched to a running server according to the order in the cluster list. The query will be queued if a queue is available, and there is no available agent in the selected server.
- SMART.** The query will be dispatched to the fastest server according to the statistics gathered from remote servers in the cluster. Overall average response time will be minimized.
- WEIGHTED.** Weight is assigned to each server based on its relative exponentially smoothed response time. A server is then selected based on the weight. If more than one server is available, the query will not be queued even if a queue is available in the selected server.

MAX_RSP_TIME

Defines the maximum query response time that users can expect. If the value is greater than zero, Cluster Manager will attempt to achieve this limit by starting reserved servers or adjusting the configuration of remote servers if AUTO_CONFIG is on.

START_STOP_THRESHOLD

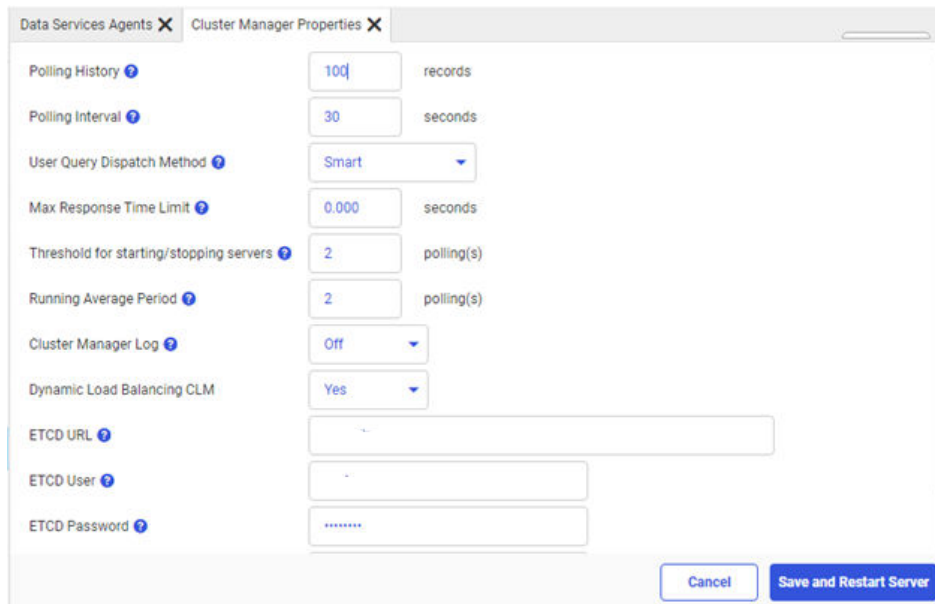
Defines the number of polls. Cluster Manager will start or stop remote reserved servers if query response time is continuously above or below MAX_RSP_TIME for the last START_STOP_THRESHOLD polls. The default value is 2.

Procedure: How to Change Cluster Manager Properties

You can set Cluster Manager properties that apply to all of your clusters.

1. Go to the *Workspace* page.
2. Right-click *Cluster Manager* on the resources pane, and click *Properties*.

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



The screenshot shows the 'Cluster Manager Properties' dialog box. It contains the following fields and controls:

- Polling History:** Input field with '100' and 'records' label.
- Polling Interval:** Input field with '30' and 'seconds' label.
- User Query Dispatch Method:** Dropdown menu with 'Smart' selected.
- Max Response Time Limit:** Input field with '0.000' and 'seconds' label.
- Threshold for starting/stopping servers:** Input field with '2' and 'polling(s)' label.
- Running Average Period:** Input field with '2' and 'polling(s)' label.
- Cluster Manager Log:** Dropdown menu with 'Off' selected.
- Dynamic Load Balancing CLM:** Dropdown menu with 'Yes' selected.
- ETCD URL:** Text input field.
- ETCD User:** Text input field.
- ETCD Password:** Password input field with masked characters.

At the bottom right, there are two buttons: 'Cancel' and 'Save and Restart Server'.

3. Make any changes needed to the properties.
4. Click *Save and Restart CLM*.

Reference: Cluster Manager Properties

With the exception of UDP Port Number, changing Cluster Manager settings will require restarting the Cluster Manager.

UDP Port Number

Defines the UDP port number that the Cluster Manager is listening on for alternate node queries from applications.

Polling History

Defines how many statistics are kept for each remote Reporting Server. The default value is 100.

Polling Interval

Defines how often the Cluster Manager gathers statistics from remote Reporting Servers in seconds. The default value is 30.

User Query Dispatch Method

Defines how Cluster Manager dispatches user queries to remote Reporting Servers. The options are:

- Adhere to the settings in Cluster Manager.** The cluster will use the option set in the Cluster Manager settings. This is the default value.
- RANDOM.** The query will be dispatched randomly to a running Reporting Server in the cluster. Every Reporting Server will have the same weight. The query will be queued if a queue is available, and there is no available agent in the selected server.
- ROUND-ROBIN.** The query will be dispatched to a running Reporting Server according to the order in the cluster list. The query will be queued if a queue is available, and there is no available agent in the selected Reporting Server.
- SMART.** The query will be dispatched to the fastest Reporting Server according to the statistics gathered from remote Reporting Servers in the cluster. Overall average response time will be minimized.
- WEIGHTED.** Weight is assigned to each Reporting Server based on its relative exponentially smoothed response time. A Reporting Server is then selected based on the weight. If more than one Reporting Server is available, the query will not be queued even if a queue is available in the selected Reporting Server.

Max Response Time Limit

Defines the maximum query response time that users expect. If the value is greater than zero, Cluster Manager will attempt to achieve this limit by starting reserved Reporting Servers or adjusting the configuration of remote Reporting Servers if AUTO_CONFIG is on.

Threshold for starting/stopping servers

Defines the maximum query response time that users expect. If the value is greater than zero, Cluster Manager will attempt to achieve this limit by starting reserved Reporting Servers or adjusting the configuration of remote Reporting Servers if AUTO_CONFIG is on.

Running Average Period

Defines the number of pollings used in averaging statistics. The default value is 2.

Cluster Manager Log

Determines whether Cluster Manager will log cluster statistics in the EDACONF/clmprint.log file.

The maximum number of lines for each cluster node is defined by the POLLING HISTORY parameter in the odin.cfg file. Each line contains the following fields, which are separated by commas, in the following order:

- Cluster name.
- Time in yyyy/mm/dd hh:mm:ss format.
- Number of active servers.
- Total number of connections received among Reporting Servers since they started or statistics were reset.
- Total number of completed connections among Reporting Servers since they started or statistics were reset.
- Average response time since Reporting Servers started or statistics were reset.
- Total number of security failures among Reporting Servers since they started or statistics were reset.
- Total number of resource failures among Reporting Servers since they started or statistics were reset.
- Total number of FOCUS errors among Reporting Servers since they started or statistics were reset.

- ❑ Running average response time over the last RUNNING_AVG_PERIOD polling.
- ❑ Number of connections to this cluster since Cluster Manager started or statistics were reset.
- ❑ Connection arrival rate to this cluster since the first connection arrived.
- ❑ Connection completion rate since the first connection arrived.
- ❑ Running connection arrival rate over the last RUNNING_AVG_PERIOD polling.
- ❑ Running connection completion rate over the last RUNNING_AVG_PERIOD polling.

Reference: Multiple Cluster Manager Configurations

Advanced parameters have been added for multiple cluster nodes. The parameters include:

- ❑ **START_STOP_THRESHOLD.** Defines the number of polls. Cluster Manager will start or stop remote reserved Reporting Servers if query response time is continuously above or below MAX_RSP_TIME for the last START_STOP_THRESHOLD polls.
- ❑ **MAX_RSP_TIME.** Defines the maximum query response time that users expect. If the value is greater than zero, Cluster Manager will attempt to achieve this limit by starting reserved Reporting Servers or adjusting the configuration of remote Reporting Servers if AUTO_CONFIG is on.
- ❑ **DISPATCH_METHOD.** Defines how Cluster Manager dispatches user queries to remote Reporting Servers.

These parameters work on the cluster node level and override the values defined in the Cluster Manager node.

Cluster Manager Global Monitoring Console

The Cluster Manager Global Server Monitoring Console monitors Reporting Servers, regardless of whether the server is explicitly added to the Cluster Manager Server or registered to the Cluster Manager Server using the Auto Cluster Manager Registration feature. This console includes monitoring capabilities to identify the up or down status of servers, average query response time, and the number of concurrent tasks running. You can also use it to perform various administrative tasks.

To access the Cluster Manager Global Server Monitoring Console in order to see the remote servers that are defined, navigate to the Workspace page of the Reporting Server browser interface for the Cluster Manager Server, right-click *Cluster Manager*, and click *Remote Servers*. The list of cluster servers and dynamic servers opens, as shown in the following image.

Type	Status	Cluster/Server	Agents Now	Agents Peak	Connections Total	Connections Completed	Avg. Response Time	Running Avg. Response Time	Queued Now
	-	Default Cluster	4	2	4	4	0.109 seconds	0.000 seconds	0
	running	Inxx64r8.lbi.com.10901	2	1	2	2	0.102 seconds	0.000 seconds	0
	running	Inxx64r6.lbi.com.10901	2	1	2	2	0.117 seconds	0.000 seconds	0
	-	CLUSTER2	4	2	4	4	0.109 seconds	0.000 seconds	0
	running	LNXR6	2	1	2	2	0.117 seconds	0.000 seconds	0
	running	LNXR8	2	1	2	2	0.102 seconds	0.000 seconds	0

In this image, the two Reporting Servers that are part of CLUSTER2 were explicitly configured remote Reporting Servers. They display with black server icons. Reporting Servers that were auto-registered display with green server icons, and are part of the Default Cluster group.

You can also use the Cluster Manager Global Server Monitoring Console to access remote server agent information and statistics and to control (kill) remote agents, as well as turn traces on or off for individual remote Reporting Servers. Right-click the *Default Cluster* group or any other cluster name, and click *Remote Data Service Agents*, as shown in the following image.

Type	Status	Cluster/Server	Agents Now	Agents Peak	Connectio
	-	Default			
	running	Inxx64r8			
	running	Inxx64r6			
	-	CLUSTER2	4	2	
	running	LNXR6	2	1	
	running	LNXR8	2	1	

Statistics


Remote Data Service Agents

The agent list for all reporting servers in all clusters will appear, as shown in the following image.

<input type="checkbox"/>	Remote Server ↓	Tscomid ↓	Service ↓	State ↓	Command ↓
<input type="checkbox"/>	LNXR6	2	DEFAULT	idle	---
<input type="checkbox"/>	LNXR6	3	WC_DEFAULT	idle	TABLEF
<input type="checkbox"/>	LNXR6	4	WC_DEFAULT	idle	TABLEF
<input type="checkbox"/>	LNXR6	5	SCHED_DEFAULT	idle	?
<input type="checkbox"/>	LNXR6	6	SCHED_DEFAULT	idle	TABLEF

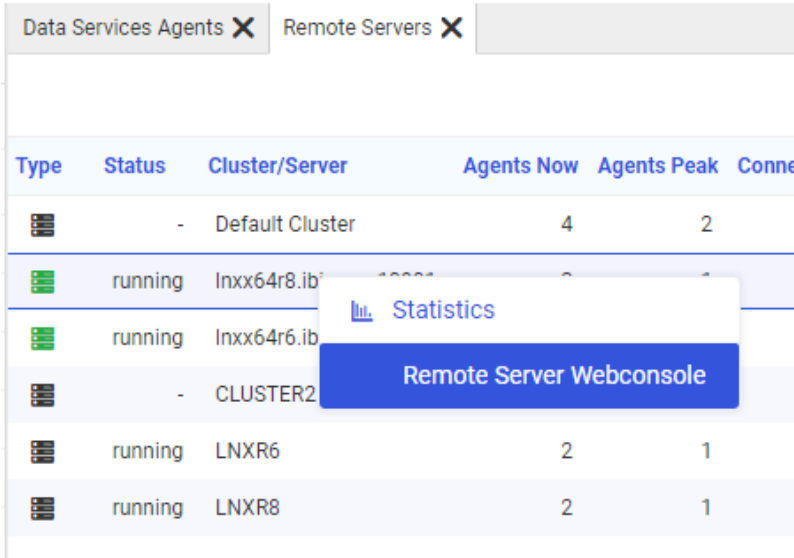
From this page, you can right-click to monitor Statistics for the agents, or to Save Last FOCUS Command, which allows you to view and save the last command run by the agent, as shown in the following image.

<input type="checkbox"/>	Remote Server ↓	Tscomid ↓	Service ↓	State ↓
<input type="checkbox"/>	LNXR6	1	DEFAULT	idle
<input type="checkbox"/>	LNXR6	2	WC_DEFAULT	idle
<input type="checkbox"/>	LNXR6	3	SCHED_DEFAULT	idle

To kill a particular agent or group of agents, select the check marks for the desired agents, and click the  icon.

To enable or disable traces on a particular remote reporting server, use the  icon.

The Cluster Manager Global Server Monitoring Console does not allow direct viewing of remote Reporting Server traces, but you can use the monitoring console to access the sign-in page of the Reporting Server browser interface for a remote Reporting Server. From the console, you can view traces, as well as perform other typical Reporting Server browser interface activities. To access the Reporting Server browser interface of a remote Reporting Server, on the Cluster Manager Global Server Monitoring Console page, right-click the remote server node name, and click *Remote Server Reporting Server browser interface*, as shown in the following image.

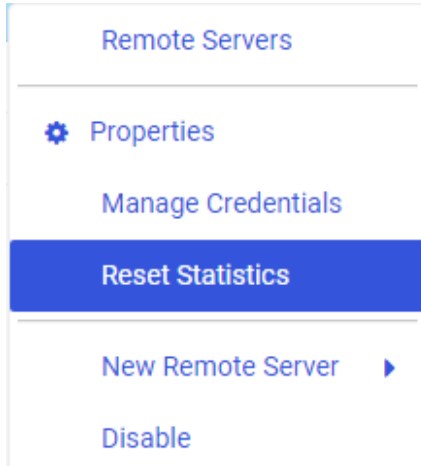


If a remote server is stopped or becomes unavailable, the Global Monitoring Console will reflect the change in the status column, as shown in the following image.

Type	Status	Cluster/Server	Agents Now
	-	CLUSTER2	4
	disconnected	LNXR6	
	running	LNXR8	2
	disconnected	Inxx64r6.ibi.com.10901	
	running	Inxx64r8.ibi.com.10901	2

Resetting Cluster Manager Statistics

The Cluster Manager Server keeps a statistical log of the global server activities. To view it, navigate to the Workspace page, right-click *Cluster Manager*, and click *Remote Servers*. To reset the log, right-click *Cluster Manager*, and click *Reset Statistics*, as shown in the following image.



Chapter 9

Listeners and Special Services

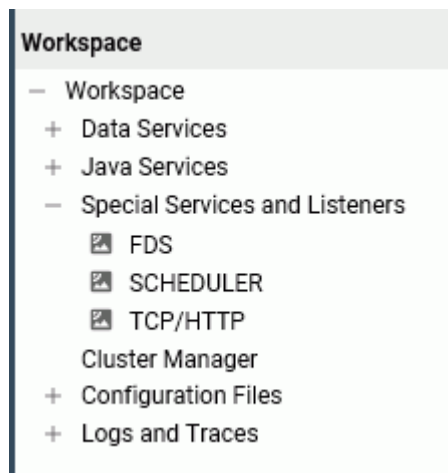
The following section describes how to manage listeners and special services.

In this chapter:

- [Using the Listeners and Special Services Configuration Panes](#)
- [Special Services](#)
- [Listeners](#)

Using the Listeners and Special Services Configuration Panes

The server provides special services and listeners, which are available from the Special Services and Listeners folder of the Workspace resources tree.



In most instances, there is no need to alter the default configurations for these features.

Special Services

The special services are:

- FDS.** FOCUS Database Server.
- SCHEDULER.** Used to schedule flows and stored procedures, and to manage deferred requests and reports.

- ❑ **Node Package Manager and Node.js In-Document Analytics Caching.** As of Release 8207.28, the Node.js Caching feature enables Reporting Server independent caching, which allows external access to Reporting Server In-Document Analytics outputs. This feature is normally configured automatically at server installation time, if Node Package Manager (npm) and Node.js are detected during the installation. This feature does not currently have control features on the Reporting Server browser interface. For more information on how to test, configure, or adjust this feature, see [Caching Support for Node.js](#) on page 835.

Running and Configuring the FOCUS Database and FOCUS Database Server

The server supports its own internal database known as a FOCUS Database, and includes its own FOCUS Database Server, also called a Sink Machine, for multi-user use. Sites running z/OS have the option of using the integrated FOCUS Database Server or keeping their FOCUS files on a legacy FOCUS Database Server, which runs as a separate batch job.

In many simple cases, syntax, such as ON TABLE HOLD AS MYAPP/MYDATA FORMAT FOCUS, will generate all that is needed for proper initial creation internally. However, more precise creations require a USE statement. On z/OS, a DYNAM or a JCL DD is also required, and the USE syntax will vary if a legacy database server is in use.

The general syntax of the USE statement is:

```
USE
XXX [NEW]
ZZZ [ON server]
XXX01 AS XXX [ON server] READ
XXX02 AS XXX [ON server] READ
{path}AAA.foc AS AAA [ON server]
{path}BBB.foc AS BBB NEW
{app}/CCC AS CCC [ON server]
{app}/AAA AS AAA NEW
END
```

The AS phrase effectively allows you to access a file under an alternate name, with a different Master File or synonym, but it also allows you to concatenate the access of data partitioned as separate files. The AS phrase is also required syntax for path and app use. Syntax using full-path, native file names is only applicable to non-z/OS platforms (z/OS uses DYNAMs or DD names to point at specific datasets).

The NEW parameter indicates that the file is being specified for creation purposes (CREATE FILE). Creates may not be done on a database server. They must be done independent of the server and then placed under server access.

The ON phrase indicates the use of a database server. On z/OS, the choices for server names are FOCUSU01 (server integrated) or FOC SBS (FOCUS legacy). All other platforms use the server name FOCUSU01. The FOCUSU01 node is a pre-configured and reserved name for the integrated FOCUS Database Server and may not be used for any other purposes. The server communication file (ODIN), can be configured to support additional FOCUS Database Server nodes and/or have additional nodes that point to the nodes of FOCUS Database servers configured under other servers.

When using an AS phrase in conjunction with an ON phrase for files where the physical name (not including path or extension) and the AS name do not match, the READ parameter is required to prevent an unintended opening for write.

Important: On z/OS, a single file cannot be accessed by both a legacy and an integrated FOCUS Database Server. You must select one or the other to manage access or you will encounter queue conflicts.

Example: **Using the Integrated FOCUS Database Server From a Global Server Profile (suprof.prf or edasprof.prf)**

Using the Reporting Server browser interface, edit either suprof.prf (preferred) or edasprof.prf (alternative). On z/OS, add DYNAM ALLOC allocations and USE statements. On other platforms, only add USE statements.

For example, use the following syntax to enable the FOCUS Database Server to control access to the CAR data source.

On z/OS:

```
DYNAM ALLOC FI CAR DA dsname SHR REU
USE
CAR ON FOCUSU01
END
```

On other platforms:

```
USE
CAR ON FOCUSU01
END
```

Example: **Using the Integrated FOCUS Database Server With a JCL Statement (z/OS Only)**

IRUNJCL JCL allocation

Add JCL statements for the FOCUS files to be managed, for example:

```
//CAR DD DISP=SHR,DSN=dsname
```

The following sample USE command enables the FOCUS Database Server to control access to the CAR data source:

```
USE
CAR ON FOCSU01
END
```

Example: **Using the Integrated FOCUS Database Server From Other Profiles or in a Procedure (FOCEXEC)**

The syntax is similar when used in a profile, but on z/OS, the DYNAM uses the PERM option and is equivalent to a JCL allocation using IRUNJCL JCL.

For example, use the following syntax to enable the FOCUS Database Server to control access to the CAR data source:

On z/OS:

```
DYNAM ALLOC FI CAR DA dsname SHR PERM
USE
CAR ON FOCSU01
END
```

On other platforms:

```
USE
CAR ON FOCSU01
END
```

Example: **Using the Legacy FOCUS Database Server (Separated Batch Job, z/OS Only)**

The following sample code identifies a legacy FOCUS Database Server using the odin.cfg node block FOCSBS:

```
USE
CAR ON FOCSBS
END
```

This sample node block resides in the odin.cfg file and can be constructed using the Reporting Server browser interface:

```
NODE=FOCSBS <==== USS server (odin.cfg file)
BEGIN
  SUBSYS = xyzw <==== optional parameter to specify IBI subsystem name
                    (not needed if you are using the name IBIS)
  PROTOCOL=SBS
  CLASS=SUCLIENT
  PORT=X.Y.Z <==== communications dataset
END
```

The Scheduler and Deferred Service


You can configure the settings that are used to manage the Scheduler and deferred requests and reports from the Reporting Server browser interface Scheduler Configuration pane. The parameters are stored in the `edaserve.cfg` file. You can also view and configure the Scheduler Agents page.

Procedure: How to Display and Customize the Scheduler Agents Page

1. Go to the *Workspace* page.
2. Click *Monitor*, then click *Scheduler Agents*, or expand *Special Services and Listeners* on the resources pane, right-click *SCHEDULER*, and click *Scheduler Agents*.

The Scheduler Agents page opens.

3. To customize the columns that appear on the page, click *Choose Columns* on the menu.
Select the check boxes for the columns you want to appear, as shown in the following image.



4. Click *OK*.
5. To customize the states that appear on the page, click *Choose States* on the menu bar.
All states appear by default. Clear the check boxes for the states you want to remove.
6. Click *OK*.
7. To activate or deactivate a Scheduler agent, select the check box and click *Activate Selected* or *Deactivate Selected*.

Reference: Scheduler Agent Columns

You can add the following columns to the Scheduler Agents page.

ID

Identification number associated with the scheduler agent.

State

The current state of the scheduler agent (Starting, Scheduled, Running, Failed, Killed, Killing, Quiesced, Quiescing, Inactive, Reloading, or Done).

Application

Application name of the scheduled procedure.

Procedure

Name of the scheduled procedure.

Report

Name of the scheduled report.

User ID

User credentials used to run the scheduled procedure.

Status

The current status for this procedure (Active or Inactive).

Schedule Type

The schedule type for this procedure (Run Once, Recurring, Multi-Day, Run when server starts, CMASAP, CMASAP Iterator, or CMASAP Iterator Group).

Start Date/Time

The initial start date and time for the schedule.

End Date/Time

The date and time on which this schedule expires.

Intraday Start

The start time of the intraday restriction interval.

Intraday End

The end time of the intraday restriction interval.

Interval

For recurring schedules, the number of minutes, hours, days, weeks, months, or years between runs.

Days of week

The days of the week on which the schedule will run.

Days of month

The days of the month on which the schedule will run.

E-Mail on Failure

The email address where confirmation will be sent on failure of each scheduled run.

E-Mail on start

The email address where confirmation will be sent on start of each scheduled run.

E-Mail on completion

The email address where confirmation will be sent on completion of each scheduled run.

Last Run Time

The time of the last scheduled run.

Next Run Time

The time of the next scheduled run.

Tscomid

The identification number for the data service agent running the scheduled procedure.

Job ID

The job ID.

Parent Job ID

The parent job ID.

Reference: Scheduler Agent States

By default, the Scheduler Agents includes all states. You can remove the following states to limit what is displayed on the page.

Main states

SCHEDULED

The agent is waiting for its next scheduled run time. When it arrives, the agent attempts to establish a connection with a tscom3 agent to execute the scheduled procedure. When a connection is established, the agent enters the RUNNING state. The agent may transition into the FAILED state if a problem occurs. The administrator may set the agent to transition into the QUIESCED state by issuing a 'Quiesce Agent' request.

INACTIVE

An idle state. The procedure has been marked inactive. No action is required.

DONE

An idle state. The agent will not run because the scheduled end time has passed. The operation is complete and no further action will be initiated by the agent.

RUNNING

A connection has been established with a tscom3 agent to execute the procedure, which has begun to execute. Attempting to kill this agent will kill the underlying tscom3 agent process and cause a transition to the KILLED state. A critical error, such as a lost or broken connection, will move the agent into the FAILED state.

FAILED

A critical error occurred while the agent was trying to connect to a tscom3 agent or while executing the procedure. Administrator action is required to move the agent out of this state. Removing a FAILED agent forces a re-scan and creates a new agent for the procedure.

QUIESCED

The agent has been placed in this state by an administrator request to stop a SCHEDULED agent. No connection exists between this agent and a tscom3 agent. Without administrator action, the agent will remain in this state permanently. To transition from this state, the administrator must remove the agent. This forces a re-scan and creates a new agent for the procedure.

KILLED

The agent has been placed in this state by an administrator request to kill a RUNNING agent. The underlying tscom3 process has been killed. No active connection exists between this agent and tscom3. Without administrator action, the agent will remain in this state permanently. To transition from this state, the administrator must remove the agent. This forces a re-scan and creates a new agent for the procedure.

Transitional states**STARTING**

The agent should successfully pass through this state into a SCHEDULED, DONE, INACTIVE, or FAILED state, with no intervention.

QUIESCING

The agent will transition from this state into a QUIESCED state, with no intervention.

KILLING

The agent should successfully pass through this into a KILLED state, with no intervention.

RELOADING

A RELOAD AGENT request was received to reload a QUIESCED, DONE, FAILED, KILLED, or INACTIVE agent. The agent should successfully pass through this state into a SCHEDULED, DONE, INACTIVE, or FAILED state, with no intervention.

RESTARTING

The scheduled agent has failed, but is automatically restarting because of the sched_restart_failed option. The agent transitions into a RUNNING state after sched_restart_interval seconds.

Procedure: How to Set Scheduler and Deferred Management Properties

1. Go to the *Workspace* page.
2. On the resources pane, expand *Special Services and Listeners*.
3. Right-click *SCHEDULER*, and click *Properties*.

The Scheduler Configuration page opens. The page has *Scanning*, *Execution*, *Logging and Output*, and *E-Mail Notification* sections.

4. To set properties for the Scheduler or for deferred execution of requests, enter values in the corresponding fields.
5. Click *Save and Restart Scheduler*.

Reference: Scheduler Scanning Properties

The Scheduler Scanning properties are:

`sched_autostart`

If set to *y*, scanning for scheduled events is started when a Scheduler service starts. If a service is not running, turning on `sched_autostart` will not make scanning available.

If set to *n*, a user may still start the scanning manually using the *Force Scheduler Scan* option.

Depending on the number of files in the application path, `sched_autostart` may affect server performance.

The default value is *n*.

`dfm_autostart`

If set to *y* (the internal default), the initial scan for dfm jobs is performed with the start of DFM listener/Scheduler Service. If it is set to *n*, the user may still initiate a scan manually. Turning the scanning on may affect the server performance on startup, depending on the number of files in the dfm directory.

`dfm_app`

Defines the application name where deferred requests and responses are stored.

`dfm_dir`

Defines the location where deferred requests and responses are stored. This is ignored if `dfm_app` is set.

Reference: Scheduler Execution Properties

The Scheduler Execution properties are:

`sched_service`

defines the agent service the Data Migrator Scheduler will use when running DM flows. If this keyword is not set, the DEFAULT service will be used.

`sched_restart_interval`

Specifies the time interval in seconds between restart attempts for all restartable flows. The default value is 60 seconds.

`sched_restart_failed`

Determines whether scheduled agents will be restarted after they fail. The valid values are:
n - Failed agents will be restarted at most *n* times.

0 - Failed agents will not be restarted.

-1 - Failed agents will be restarted infinitely.

The default value is -1.

`CMASAP_clear_interval`

Specifies how often (in seconds) CMASAP agents are removed from the agents list. The default value is -1, meaning never.

Reference: Scheduler Logging and Output Properties

The Scheduler Logging and Output properties are:

`sched_log_lines`

Defines the maximum number of lines per request that the Scheduler will write to the log. When this number of lines is reached, logging stops. Exceeding this limit for one request does not affect logging for other requests.

When set to -1, the number of lines is unlimited. When set to 0, logging is disabled. The default value is -1.

`sched_standalone_log`

If checked, generates a separate log file for each job.

`sched_dmlog_app`

Is the name of the application in which the scheduler separate log and statistics files are located.

`sched_log_commit_interval`

Specifies the maximum time interval in seconds that can elapse before log data is committed to the log. Used in conjunction with `sched_log_commit_maxlines`, these settings can fine-tune Scheduler logging behavior. The logging subsystem will commit data to the log whenever one of these maximums has been reached. Finding the best commit interval is a compromise among latency (the delay between the time a log message is received and when it is committed to the log database), reliability, and performance. Longer intervals enable greater I/O efficiency, but increase the potential amount of log information lost in a crash. The default value is 10.

`sched_log_commit_maxlines`

Specifies the maximum number of log lines to collect before committing them to the log. Used in conjunction with `sched_log_commit_interval`, these settings can fine-tune Scheduler logging behavior. The logging subsystem will commit data to the log whenever one of these maximums has been reached. Larger numbers enable greater I/O efficiency, but increase the potential amount of log information lost in a crash. The default value is 1000.

`Merge formatted output into log`

Changes value of `sched_log_output_destination` keyword.

When checked, it specifies that only the job Log is stored in the Scheduler log (`sched_log_output_destination=Log`). The output is stored in the `dfm_dir` location.

When unchecked, it specifies that the job Log and Output are stored in the Scheduler log (`sched_log_output_destination=LogAndOutput`).

`dfm_maxage`

Defines the maximum number of days that deferred reports are kept in the server after they are created. The value zero (0) means the age is unlimited. The default value is 30.

`dfm_maxoutput`

Defines the maximum size in kilobytes of a deferred report. Reports that are over this limit will be removed.

On z/OS PDS deployment, the number represents a number of records. The valid value is any number between 0 and 15,728,688 inclusive.

On all other platforms, the number represents a number of kilobytes or megabytes with the suffix K or M. A valid value is any number between 0 and 65535 inclusive.

The default value is zero (0), which means the output is unlimited.

Reference: Scheduler E-Mail Notification Properties

The Scheduler E-Mail Notification properties are:

`On Start`

Sends email when the procedure starts.

`On Completion`

Sends email when the procedure completes.

`On Failure`

Sends email if the procedure fails.

`sched_max_att`

Enter the e-mail attachment size limit in Kb.

`sched_email_address`

Defines the list of email recipients. The list can contain one or multiple email addresses and/or user IDs. Multiple entries should be separated by semicolons (;).

`dfm_email`

When checked, the scheduler Email Notification settings are applied to deferred requests. The default is unchecked, which does not apply the scheduler email settings to deferred requests.

`dfm_email_from_userid`


When the response of a deferred request is ready to be retrieved, an email notification can be sent to a specific email address. If this keyword is set to n, no email notification will be sent. If it is set to y and there is no email address specified in the deferred request, an email notification will be sent to the email address of the requestor specified in the admin.cfg configuration file or in the corresponding security provider. The default value is n.

***Procedure:* How to Manage Tenants**

This option allows you to manage Scheduler User IDs used by the Scheduler for scanning (sched_scan_id) and running scheduled jobs (sched_run_id).


1. In the Special Services and Listeners folder of the Workspace page, right-click *Scheduler* and click *Manage Tenants*.

The *Manage Scheduler Userids* page opens, as shown in the following image.



2. To add a new user or run ID, click *New* on the menu bar.

The New Tenant dialog box opens, as shown in the following image.



Select values for the following parameters.

ScanID

Determines what user ID the Data Migrator scheduler uses to scan the application path for scheduled Data Migrator flows. This can be used to restrict the set of application directories scanned by the scheduler to a subset of the APP PATH of the effective administrator. Select a user from the drop-down list. This sets sched_scan_id.

Type

Can be one of the following.

- <server_admin_id>**. This is the default value.

Setting sched_run_id to <server_admin_id> uses the first valid Server Administrator user ID defined in the admin.cfg file to run the Data Migrator job by the Scheduler. Usually this would be the effective Server Administrator.

- <user>**.

Setting sched_run_id to <user> uses the user ID from the Data Migrator job to run the job by the Scheduler.

- RunID.** If you select this option, a RunID field opens. You must select a user from this list. This value determines what user ID a Data Migrator flow uses to run scheduled flows, either `server_admin_id` or `user`.

Note: Data Migrator flows submitted from the Data Migrator desktop interface, the Reporting Server browser interface or CMRUN are run under the userid that submits the flow.

3. Click *OK*.

The ID is added to the *Manage Scheduler Userids* page.

The pairs of `sched_scan_id` and `sched_run_id` are written to the `edaserve.cfg` file as follows:

```
sched_scan_id = PTH\srvadmin
BEGIN
  sched_run_id_type = server_admin_id
END
```

or

```
sched_scan_id = PTH\scanid1
BEGIN
  sched_run_id_type = user
END
```

or

```
sched_scan_id = PTH\scanid2
BEGIN
  sched_run_id_type = PTH\runid1
END
```

Multiple `sched_scan_id` and `sched_run_id` pairs or tenants can be defined. During scanning, the Scheduler will use each tenant `sched_scan_id` to scan and update the Scheduler.

If multiple tenants have the same application in the `sched_scan_id` APPPATH, the jobs from this application will be scheduled only once, by the first tenant used during scanning.

4. To delete a tenant, select one from the list and click *Delete*.
5. To edit a tenant, select one from the list and click *Properties*.
6. When you are finished, click *OK*.

Procedure: How to View Deferred Process Statistics

1. Go to the *Workspace* page.

2. On the resources pane, expand *Special Services and Listeners*.
3. Right-click *SCHEDULER*, and click *Statistics*.

The Statistics SCHEDULER page displays the following information in the *DFM* section:

- DFM_DIR Available Disk Space (KB)
- Number of Requests Done Since Startup
- Number of Response Ready

For more information, see [Statistics for an Individual Special Service](#) on page 691.

Reference: Extensions for Deferred Files

The following table lists possible extensions for the deferred files listed in the *dfm_dir* directory.

Extension	Description
RQO	Request file.
RQD	Data file. Contains user ID, optional flags, and so on.
RQP	Request is being executed.
RQF	Request is waiting to be executed.
DEL	Request is deleted.
RPO	Response file. Contains the whole report for one request.
RPF	Response is ready.
RPE	Response exceeds maximum limit.

Forcing a Scheduler Scan

A scheduler scan is initiated on server startup, if `sched_autostart=y`. In addition, the scheduler is updated for a procedure if you edit schedule information for a procedure or Data Migrator job using *Manage* on the Schedule and Email menu from the Reporting Server browser interface or Data Migrator desktop interface. At other times, you can force scanning of scheduler and deferred jobs in the following ways.

- ❑ On the Applications page, right-click an application folder, point to *Schedule and E-Mail*, and click *Force Scheduler Scan*.
- ❑ On the Workspace page, open Special Services and Listeners, right-click Scheduler, and click *Force Scheduler Scan*.
- ❑ Issue the following command in the server command window.

```
edastart -forcscan
```

Storing Deferred Reports in an Application Directory or an SQL Repository

Deferred requests and reports can be written to any application, including applications that are stored in an SQL repository.

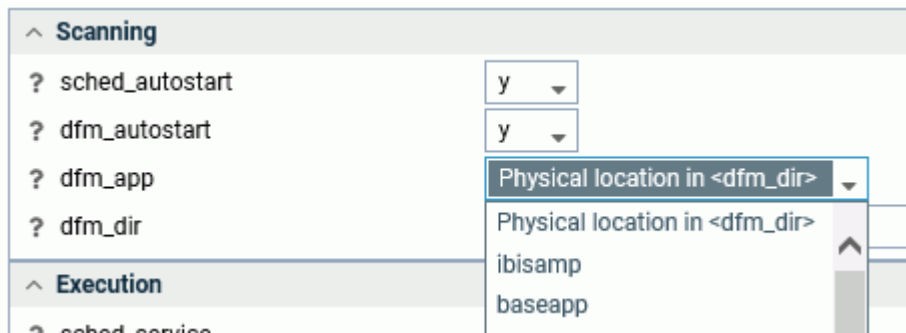
To route deferred requests to any application:

1. Sign on to the server with a server administrator ID.
2. On the Workspace page, expand *Special Services and Listeners*.
3. Right-click *SCHEDULER*, and click *Properties* from the shortcut menu.

The Scheduler Configuration page opens.

By default, `dfm_app` is set to a physical directory location.

4. You can accept or edit the `dfm_dir` location or select an application from the `dfm_app` drop-down list, as shown in the following image.



If you select an application, it will override the directory identified in `dfm_dir`, and the `dfm_dir` text box will be hidden.

5. Click *Save and Restart Scheduler*.

Deferred requests and output will be returned to that application.

Listeners

The listeners are:

- HTTP
- TCP
- SOAP

These listen for activity on their respective ports.

Note: To include a SOAP listener, you need to configure one. For more information, see [How to Configure a SOAP Listener](#) on page 752.

Configuring Listeners

All listeners have the same type of basic parameters. They differ in the additional parameters that are available. All three types have Security parameters. HTTP also includes Sessions Control, Aliases and Miscellaneous Settings.

To Configure a Listener:

1. Go to the *Workspace* page.
2. Open *Special Services and Listeners* on the resources pane.
3. Right-click a listener and click *Properties*.

The Listener Configuration page opens.

4. In the *Basic* section, you can enter a value for the *HOST* parameter.
5. If you need to set additional parameters, expand the corresponding section and enter values in the parameter fields.
6. Click *Save and Restart Server*.

Reference: Listener Basic Parameters

The HTTP, TCP, and SOAP listeners have the same type of basic parameters.

NODE

Defines the logical name of a node block.

The settings are:

For HTTP - LST_HTTP.

For TCP - LST_TCP.

For SOAP - LST_SOAP.

PORT

Defines the port number that a listener is listening on. The default value for the HTTP listener is 8117. The default value for the TCP listener is 8116.

HOST

Defines an IP address that a listener is listening on. If blank, the listener will listen on all IP addresses.

Reference: Additional Listener Parameters

The HTTP, TCP, and SOAP listeners have the following security parameter:

RESTRICT_TO_IP

Defines the name of a host or IP address(es) that will be accepted by the listener. The syntax is

```
hostname, xxx.xxx.xxx.xxx, yyy.yyy.yyy, ...
```

The address must be in base 256 standard dot notation. For platforms that support IPv6, IPv6 notation can be used if it is enabled (for example, 2001:1b1:719:1b1:203:baff:fe0a:fe23)

The internal default is *.*.*.*, which allows all IP addresses.

Note: You can use a wildcard to mask an entire section of the address, as in the following examples:

```
172.204.201.*
172.*.*.*
172.204.*.*
2001:1b1:719:1b1:203:baff:*:fe23
```

Partial masking is not supported, for example:

```
172.204.201.1*
172.204.201.*23
```

The HTTP listener has the following additional parameters:

Security

SECURITY

Defines the authentication protocol used by a secured listener. This parameter can only be set if OPSYS is the primary security provider. The valid values are:

- Default. This is the default value, which does not add a SECURITY attribute in the HTTP listener node of the odin.cfg file.
- Kerberos or NTLM. This setting adds the attribute SECURITY=KERBEROS in the HTTP listener node of the odin.cfg file.

Kerberos is only available on Windows, Linux, and IBM i. Kerberos supports passing an authentication ticket between multiple computers (multiple hops). For example, the Kerberos authentication ticket can be passed from a browser computer to a server computer to a DBMS computer. For Kerberos authentication to work, computers must be configured for Kerberos. Contact your administrator for instructions on configuring Kerberos in your environment.

NTLM is only available on Windows and supports Integrated Windows Authentication (IWA) between two computers (for example, a browser computer to a server computer or a server computer to a DBMS computer).

- NTLM only. This setting adds the attribute SECURITY=IWA (Integrated Windows Authentication) in the HTTP listener node of the odin.cfg file.

IWA is only available on Windows and supports authentication between two computers (for example, a browser computer to a server computer or a server computer to a DBMS computer).

To configure the server for Kerberos on Linux and IBM i, you must set the following properties:

1. First, on the Access Control page, enter the SPN (Service Principal Name) in the krb5_srv_principal field. Contact your Domain Administrator for information.
2. Select *Kerberos or NTLM* for the SECURITY parameter in the HTTP Listener properties.
3. Create local user IDs on Linux or IBM i matching domain users that will connect to the server using Kerberos.

The following table shows the effects of the Security settings on authentication.

SECURITY Setting	SECURITY attribute in odin.cfg	Server Machine Configured for Delegation Browser and Server on Same Machine	Server Machine Configured for Delegation Browser and Server on Different Machines	Server Machine Not Configured for Delegation Browser and Server on Same Machine	Server Machine Not Configured for Delegation Browser and Server on Different Machines
Default	None	Explicit Connection. User prompted for credentials.	Explicit Connection. User prompted for credentials.	Explicit Connection. User prompted for credentials.	Explicit Connection. User prompted for credentials.
Kerberos or NTLM	SECURITY = KERBEROS	IWA	KERBEROS	IWA	Fails to connect and issues Session Lost message. Edaprint contains message that the server was unable to initialize Kerberos and an explicit connection is required.
NTLM only	SECURITY=IWA	IWA	IWA	IWA	IWA

LOGIN_FAILURE

Defines the type of message shown for a login error. The valid values are:

0 - Shows a general error message.

1 - Shows a precise error message.

The default value is 0.

LOGIN_HIDE_PROVIDERS

Defines whether available security providers or domains should be shown in the login screen. The valid values are:

- 0 - shows all available security providers or domains.
- 1 - hides all available security providers or domains.

CSRF_TOKEN

Defines whether a secret token should be used in all form submissions. This is used to prevent cross-site request forgery attacks. The valid values are:

- 0 - disable CSRF token.
- 1 - enable CSRF token.

SANITIZE_PARAMETERS

Defines whether the HTTP listener should sanitize all submitted form parameters. Valid values are zero (0) to disable sanitizing of form parameters, or 1 to enable sanitizing of form parameters. The default value is 1.

VERIFY_REFERERER

Defines whether to verify the HTTP referrer header. Valid values are zero (0) to disable verifying the referee header, or 1 to enable verifying the referrer header. The default value is 1.

X_FRAME_OPTIONS

Defines whether a browser should be allowed to render the Reporting Server browser interface in a frame. Valid values are:

- NONE.** This is available for the HTTP Listener only. The Reporting Server browser interface can be displayed in a frame.
- DENY.** The Reporting Server browser interface cannot be displayed in a frame.
- SAMEORIGIN.** The Reporting Server browser interface can only be displayed in a frame on the same origin as the page itself. This is the default value.

Enable HTTPS

Enables HTTPS support, either for OpenSSL or Microsoft Windows SSL.

Both types have the following additional parameters:

`SSL_CERTIFICATE`

Defines the SSL certificate. The default certificate name for Microsoft SSL is `ipaycert.p12`.

`SSL_PASSPHRASE_E`

If the private key of the certificate is encrypted, an encrypted passphrase must be provided to decrypt the private key.

The following additional parameters are specific to OpenSSL:

`SSL_PRIVATE_KEY`

Defines the file that contains the private key for the listener. It must correspond to the public key embedded in the PEM certificate and must be in PEM format.

`SSL_CA_CERTIFICATE`

Defines the file containing the trusted CA certificate in PEM format. It is used to verify the client certificate. If the client fails to send the certificate or if verification fails, connections will be rejected. More than one CA certificate may be present in the file.

The following additional parameter is specific to Microsoft SSL:

`SSL_FRIENDLY_NAME`

Indicates a name used to identify the certificate in the PKCS#12 file, if the file contains more than one certificate.

`SSL_AUTH_CLIENT`

Defines whether client authentication should be enabled. Valid values are:

- 0 (zero)**. The client is not asked for a certificate. There is no client authentication. This is the default value.
- 1**. The server requests a certificate from the client. All received certificates are validated. If the client fails to send a certificate or verification fails, connections will be rejected.

Sessions Control

`PERSISTENT_GLOBAL`

Defines whether global FOCUS variables will be persistent within a browser session. The valid values are:

0 - Global variables are not persistent within a browser session.

1 - Global variables are persistent within a browser session.

The default value is 0.

PASS_EXPIRE_NOTIFICATION

Determines whether users receive a password expiration notification by defining the number of days before expiration at which a notification will occur. The default value is 0, meaning no notification.

MAX_WEBSSESSION

Defines the maximum number of active sessions. The HTTP listener will reject connections to the Reporting Server browser interface if this number is reached. The default value is 0, meaning no limit.

TIMEOUT_WARNING

Defines how many minutes before session expiration a warning message will be displayed. The default value is zero (0), which means there is no timeout warning message.

LOG_LAST_REQUESTS

Defines the maximum number of last requests in a session that will be collected as statistics. The default value is 10.

Alias

IBI_HTML

Defines a URL alias for the primary file lookup directory. Lookup is performed first in the directory described by this alias. The default value is \$EDAHOME/etc.

HTML_HOME

Defines a comma-delimited list of directories for file lookup. Directories are searched in the order in which they appear in this list.

SESSION_EXPIRATION

This parameter has been deprecated and has been combined with the `foccache_maxage` parameter, in order to make sure that no foccache files remain when the session has expired. For more information about the `foccache_maxage` parameter, see [Application Settings](#) on page 620.

Miscellaneous Settings

DEFAULT_HOST

Defines the preferred hostname or IP address for the listener when a listener has multiple IP addresses or hostnames.

PROXY_SERVER

Runs the HTTP Listener as a proxy server.

The valid values are:

0 - Proxy server for the HTTP protocol is off.


1 - Proxy server for the HTTP protocol is on.

The default value is 0.

Procedure: How to Configure a SOAP Listener

A SOAP listener needs to be configured before it will appear in the Special Services and Listeners folder.

1. Go to the *Workspace* page.
2. Right-click *Special Services and Listeners*, point to *New*, and click *SOAP*.



The Listener Configuration page opens.

3. Enter values in the *PORT* and *HOST* fields.
4. Optionally, enter a host or IP address that will be accepted by the listener in the *RESTRICT_TO_IP* field in the Security section.
5. Click *Save and Restart Listener*.

The SOAP listener is added to the Special Services and Listeners folder.

Procedure: How to Return a JSON Answer Set Using the SOAP Listener

Using the Simple Object Access Protocol (SOAP) Listener, you can run report requests or Data Migrator flows using RESTful Web Service calls. The answer set is returned in JSON format.

1. First, configure the SOAP Listener as described in [How to Configure a SOAP Listener](#) on page 752.
2. Place the procedure to run in an application directory in the APP PATH of the Reporting Server.

Assuming you created the WebFOCUS Retail tutorial in the `ibisamp` application, the following `FOCEXEC` can be created in `ibisamp` to demonstrate the JSON Answer Set feature. If you created the WebFOCUS Retail tutorial in another application name, substitute application names where applicable.

```
TABLE FILE WF_RETAIL_LITE
SUM
  COGS_US
BY BRAND
BY PRODUCT_CATEGORY
BY PRODUCT_SUBCATEG
IF PRODUCT_CATEGORY EQ '&CATEGORY'
END
```

Note that the request has an `IF` test that requires a value for the variable `&CATEGORY`. You will supply the parameter name and value as part of the URL.

3. In a browser window, enter a URL in the following form.

```
http://host:port/rest/app/proc?param=value ...
```

where:

host

Is the Reporting Server host name.

port

Is the SOAP Listener port number.

app

Is the application that contains the procedure. This application must be on the `APP PATH` of the Reporting Server.

proc

Is the procedure name and extension.

param

Is a parameter name.

value

Is the parameter value.

Example: Returning a JSON Answer Set

For example, to run the `retail_ws.fex` procedure with the parameter `&CATEGORY=Accessories`, in the `ibisamp` application, on the host `localhost`, enter the following URL and press Enter.

```
http://localhost:8050/rest/ibisamp/retail_ws.fex?CATEGORY=Accessories
```

The following answer set is returned in JSON format. In this answer set, after the message header information and report statistics, each row is listed in an array named *rows*. The array of all rows is enclosed in square brackets ([]). Each row is enclosed in curly braces ({}) and separated from the other rows by a comma (,). The rows consist of *fieldname:value* pairs separated by commas. Alphanumeric values are enclosed in double quotation marks (").

```
{ "_ibi_Report" : {"Messages" : [ "0 NUMBER OF RECORDS IN TABLE=      10
LINES=      10"],
"rows" : [
  { "BRAND" : "Audio Technica" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      38000.00},
  { "BRAND" : "Denon" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      25970.00},
  { "BRAND" : "Grado" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      21930.00},
  { "BRAND" : "Logitech" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Universal Remote Controls" ,
    "COGS_US" :      61432.00},
  { "BRAND" : "Niles Audio" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Universal Remote Controls" ,
    "COGS_US" :      73547.00},
  { "BRAND" : "Pioneer" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      16720.00},
  { "BRAND" : "Samsung" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Charger" , "COGS_US" :      5405.00},
  { "BRAND" : "Sennheiser" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      78113.00},
  { "BRAND" : "Sony" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Charger" , "COGS_US" :      3168.00},
  { "BRAND" : "Sony" , "PRODUCT_CATEGORY" : "Accessories" ,
    "PRODUCT_SUBCATEG" : "Headphones" , "COGS_US" :      18592.00}
]
}
}
```

Other Listener Options

The Listener right-click menu also enables you to configure the servlet and refresh the WebFOCUS jar files.

Procedure: How to Configure the Servlet

1. Go to the *Workspace* page.
2. Open *Special Services and Listeners* on the resources pane.
3. Right-click *TCP/HTTP* and click *Configure Servlet*.

The Configure Servlet page opens.

4. Enter the path to your Java Development Kit in the *JDK_HOME* field or click or click the selector button (...) and navigate to it.

5. Optionally, you can elect to register the JDK by selecting *y* from the *Register JDK_HOME* drop-down menu.
6. Enter the string that uniquely identifies the servlet and is a part of the URL allowing access to the servlet in the *Context PATH* field. The default value is *ibi_apps*.
7. Click *Configure*.

Procedure: How to Refresh the TIBCO WebFOCUS Jar Files

1. Go to the *Workspace* page.
2. Open *Special Services and Listeners* on the resources pane.
3. Right-click *TCP/HTTP* and click *Refresh WF jar files*.

The *Copy server related java components into WF installation directory* page opens.

4. Enter the path to the jar files in the *IBI_Repository_Root_Directory* field or click the selector button (...) and navigate to it.
5. Click *Refresh WF jar files*.

Note: After copying the files, the application server will need to be restarted.

Troubleshooting

The Reporting Server browser interface enables administrators to access several diagnostics tools that can be used to visualize different internal information. The following section describes the pages used to access such information to perform problem analysis tasks.

In this chapter:

- [Viewing Version Information](#)
 - [Analyzing TIBCO WebFOCUS Reporting Server Activity](#)
 - [Analyzing FOCUS Database Server Activity](#)
 - [Tracing TIBCO WebFOCUS Reporting Server Activity](#)
 - [z/OS-Specific Troubleshooting for the Unified TIBCO WebFOCUS Reporting Server](#)
 - [z/OS-Specific Troubleshooting for z/OS HFS Deployment](#)
 - [Recording and Reproducing User Actions](#)
 - [Troubleshooting the Reporting Server Browser Interface](#)
 - [TIBCO WebFOCUS Reporting Server Processes](#)
 - [Gathering Diagnostic Information for Customer Support Services](#)
-

Viewing Version Information

The Version information page displays the main identification parameters of the Server configuration that was in effect when the Server was started.

On the Reporting Server browser interface menu bar, click the ? icon, then click *Version*. The Version pane opens.

The following information is displayed:

General**Configuration Date**

Indicates the configuration date of the Server.

Build date

Identifies the build.

Gen number

Identifies the gen.

Release

Identifies the release.

Source date

Identifies the source date.

Extended

WF jar seq number

WF jar release

Provides jar-related information.

Build system

Support System

Build-related and system-related information.

C Compiler

C Version

Identifies the C compiler and version.

Ajax Framework

identifies the Bindows version.

Host Name

Name of the machine where the Server is installed.

Server Name

As defined in the configuration file.

Note: Used on Windows platforms as the system service name for the Server.

Analyzing TIBCO WebFOCUS Reporting Server Activity

You can use the Reporting Server browser interface to view the current or prior server activity log (edaprint.log) from either a management perspective (Session and Connection activity) or a raw version of the log with or without filtering for certain components and IDs. By default, the Server Log chronologically records all server activity from the time the Workspace Manager was started. It contains basic activity information, as well as server start up information and IBISNAP and shutdown information. IBISNAP is a snapshot of the server environment showing various usage statistics, CPU and server times for each data agent connection, and listener and agent status. If an abend occurs, the snapshot contains additional debugger stack information that aids the Customer Support Services staff in determining the problem.

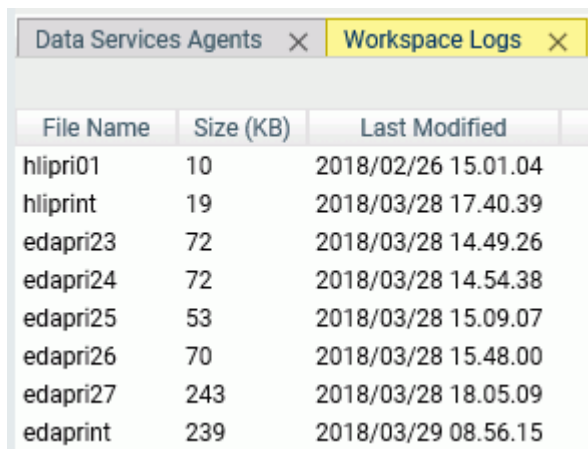
You can view the first 50,000 lines of the log file from the Reporting Server browser interface. If you wish to view lines beyond 50,000, you can open or save the file in a text editor.

By default, the server log files edaprint.log, hliprint.log, rmldata.log, and edadata.log are stored in the edatemp directory. This directory is created under the EDATEMP location which, by default, is set to EDACONF. To override this location, set the EDALOG environment variable to a physical directory prior to starting the server. The edatemp directory will still be created under EDATEMP.

Procedure: How to Access the Log Files

1. Go to the *Workspace* page.
2. Expand *Logs and Traces* on the resources tree.
3. Right-click *Workspace Logs* and click *View*.

The *Workspace Logs* page opens.

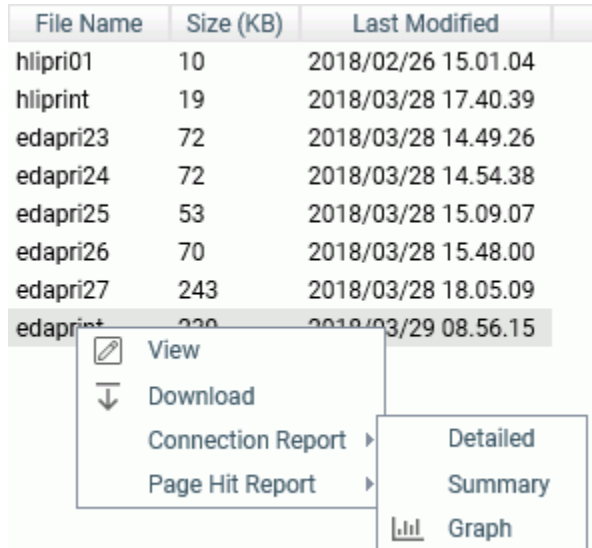


File Name	Size (KB)	Last Modified
hlipri01	10	2018/02/26 15.01.04
hliprint	19	2018/03/28 17.40.39
edapri23	72	2018/03/28 14.49.26
edapri24	72	2018/03/28 14.54.38
edapri25	53	2018/03/28 15.09.07
edapri26	70	2018/03/28 15.48.00
edapri27	243	2018/03/28 18.05.09
edaprint	239	2018/03/29 08.56.15

4. Right-click *hliprint.log* and select *View*

or

Right-click *edaprint.log* and click *View* , *Download*, or select a report type:



- View* opens the log file in a tab on the *Workspace* page.
- Download* downloads the log file to the *Downloads* folder on your PC.
- Connection Report* has *Detail*, *Summary* and *Graph* options.
- Page Hit Report* has *Detail* and *Summary* options.

The report or graph opens in a tab on the *Workspace* page.

Procedure: How to Display the Live EDAPRINT Log File

To display the live EDAPRINT log file:

1. Go to the *Workspace* page.
2. Click *Monitor*, then click *Live Console*.

The EDAPRINT log opens in a separate browser page showing connections, disconnects, and other activity happening live on the server.

Note: This option is only available when using a Firefox or Chrome browser.

Procedure: How to Open a Copy of the EDAPRINT Log File

To open a copy of the EDAPRINT log:

1. Go to the *Workspace* page.
2. Click *Diagnostics*, point to *Download*, then click *Current Workspace Log*.

Your browser will ask whether you want to open or save the file.

If you choose open, the file will open in your default text editor.

The top of the EDAPRINT log file contains all environment variables.

In addition, if JDK_HOME and JAVA_HOME are both not set, the following warning message displays in EDAPRINT:

```
W JDK_HOME= or JAVA_HOME= required for Java Listener start
```

The following shows a sample of the top of the edaprint file.

```
02/24/2016 09:57:11.753 I Workspace initialization started
02/24/2016 09:57:11.753 I EDACONF=C:\ibi\srv99\wfs
02/24/2016 09:57:11.753 I EDAHOM=C:\ibi\srv99\home
02/24/2016 09:57:11.753 I Environment Variables...
02/24/2016 09:57:11.753 I ALLUSERSPROFILE=C:\Documents and Settings\All
Users
02/24/2016 09:57:11.753 I APPDATA=C:\Documents and Settings
\se02703\Application Data
02/24/2016 09:57:11.753 I CLASSPATH=.
02/24/2016 09:57:11.769 I C:\Program Files\Java\jre6\lib\ext
\QTJava.zip
02/24/2016 09:57:11.769 I CommonProgramFiles=C:\Program Files\Common Files
02/24/2016 09:57:11.769 I COMPUTERNAME=INFORMAT-31B4D7
02/24/2016 09:57:11.769 I ComSpec=C:\WINDOWS\system32\cmd.exe
02/24/2016 09:57:11.769 I DEFLOGDIR=C:\Documents and Settings\All Users
\Application Data\McAfee\DesktopProtection
02/24/2016 09:57:11.769 I EDACONF=C:\ibi\srv99\wfs
02/24/2016 09:57:11.769 I EDAEXTSEC=OFF
02/24/2016 09:57:11.769 I EDAHOM=C:\ibi\srv99\home
02/24/2016 09:57:11.769 I EDAPRINT=C:\ibi\srv99\wfs
02/24/2016 09:57:11.769 I EDASHARE=C:\ibi\srv99\wfs
02/24/2016 09:57:11.769 I EDATEMP0=C:\ibi\srv99\wfs\edatemp
02/24/2016 09:57:11.769 I FP_NO_HOST_CHECK=NO
02/24/2016 09:57:11.769 I HOMEDRIVE=H:
02/24/2016 09:57:11.769 I HOMEPATH=\
```

Procedure: How to Set Workspace Log Properties to Control the Edaprint Log

The Workspace Log Properties page controls both the edaprint log and event routing email. For information on event routing email, see [How to Set Workspace Log Properties for Event Routing](#) on page 585.

1. Go to the *Workspace* page.
2. Open *Logs and Traces* on the resources pane.
3. Right-click *Workspace Logs* and click *Properties*.

The Workspace Logs Properties page opens, as shown in the following image.

The screenshot shows the 'Workspace Logs Properties' configuration page. It is divided into two sections: 'Event E-Mail' and 'EDAPRINT'. The 'Event E-Mail' section contains two input fields: 'email_suppress_interval' with a value of 600 and 'logging_email_edaprint' with a value of 20. The 'EDAPRINT' section contains four controls: 'edaprint_history' (input field with value 5), 'edaprint_max_lines' (input field with value 0), 'edaprint_max_days' (input field with value 0), and 'edaprint_lpname' (dropdown menu with value 'n'). At the bottom right of the form are 'Cancel' and 'Save' buttons.

4. Enter or select values for the EDAPRINT parameters.

The EDAPRINT parameters are:

`edaprint_history`

Defines the number of saved edaprint.log files to keep. Valid values are 0-99. The default value is 5.

`edaprint_max_lines`

Defines the maximum number of lines in an edaprint.log before it is archived, and a new edaprint.log is started. The default value is 0, which means unlimited.

`edaprint_max_days`

Defines the maximum number of days from the time the server is started before an edaprint.log is archived and a new edaprint.log is started. The default value is 0, which means unlimited.

Note: A new `edaprint.log` file will only be created after the `edaprint_max_days` value has been exceeded, and a connection event has occurred. For example, if the server is idle, the new `edaprint.log` file will not be created until a connection request is received.

`edaprint_ipname`

Defines whether the server should convert IP host addresses to names. It requires DNS lookups. The default value is `n`.

5. Click *Save*.

Procedure: How to Filter the Log

1. Go to the *Workspace* page.
2. Expand *Logs and Traces* on the resources pane.
3. Right-click *Workspace Logs* and click *Filter*.

The EDAPRINT Log Filter page opens.

Select an option from the Select Filtering drop-down menu.

4. If you choose *Filter by Category*, the following page opens. Set your filtering parameters and click *Set Filter*. For any item, clicking the check box without an entry in the input box filters on all items of that type. If you also enter a specific value in the input box, that filter will be executed. (For example, you can request filtering on all users, which defaults to *, or an individual user.) You can combine items from the list, mixing all items of a type with specific items, as needed.

The screenshot shows a web interface for setting filters. At the top, there are two buttons: "Set Filter" and "Clear Filter". Below them is a section titled "Select Filtering:" with a dropdown menu currently set to "Filter by Category". Underneath, there are three unchecked checkboxes: "Any Error", "Any Warning", and "Warning (Unauthorized Access or Rejected)". A note states "Use the asterisk character (*) as a wildcard". Below this are several input fields, each with a label: "User ID (user)", "Group (group)", "Role (role)", "Process ID (pid)", "TSCOM ID (tscomid)", "Session ID (sesid)", "Client Session ID (session)", and "JSCOM ID (jscomid)".

- If you choose *Filter by Timestamps*, the following page opens. Enter From timestamp and To timestamp information and click *Set Filter*.

The screenshot shows a web interface for configuring filters. At the top are two buttons: 'Set Filter' and 'Clear Filter'. Below them is a 'Select Filtering:' dropdown menu with 'Filter by Timestamps' selected. Underneath, the text 'timestamp samples:' is followed by three lines of timestamp information: '01/25/2008 10:56:47', '01/25/2008 10', and '01/25/2008'. At the bottom, there are two checkboxes. The first is labeled 'Timestamp From:' and has a text input field next to it with the label 'timestamp'. The second is labeled 'Timestamp To:' and also has a text input field next to it with the label 'timestamp'.

- Right-click *Workspace Logs* and click *View*.
The *Workspace Logs* page opens.
- Right-click *edaprint.log* and select a report type.

Reference: Controlling EDAPRINT History

The size of the server activity log, the number of files archived in the EDAPRINT history, and the maximum number of days each *edaprint.log* is maintained before it is archived and a new *edaprint* log is started are controlled by *edaprint_max_lines*, *edaprint_history*, and *edaprint_max_days*, respectively.

Limit the size of each log file by setting the maximum number of lines to a value other than zero. There is no minimum, but the recommended value is at least 1000 lines. If the limit is too low and the system is too fast (such that two archived files can have the same file system time stamp to the same second), whichever of the archived files is deleted first is undefined.

The first log file for each server start always contains more lines than the maximum because the count begins only after the *edaplog* process is created, thus it is larger by the size of the startup information. Each time the limit is reached, the current log file is archived and the new file starts with an additional identification line.

This line indicates the rank of the file in the series of continuation files, and a reference to the archived name of the first file in the series (the one with the startup information), in a format similar to the following:

```
continuation #n of edapriNN
```

The number of files kept archived is determined by setting the `edaprint_history` parameter to the maximum number of `edapriNN.log` files. Each time the server is restarted or the current log reaches maximum size, the previous `edaprint` file is moved to `edapriNN.log`, and the maximum number of files is enforced by deleting the oldest files. The effective minimum number of files kept may be as high as 2 in addition to the current `edaprint.log`, even if `edaprint_history` was set to 0 or 1 because when files are split to respect `edaprint_max_lines`, the most recent file and the most recent file containing startup information are always preserved.

Analyzing FOCUS Database Server Activity

FOCUS Database Server (FDS) is a listener that enables multiple users (or multiple agents) to update central a FOCUS database simultaneously. Database transaction events are logged to the HLIPRINT log file (`hliprint.log`), which is created in the server directory.

You can view the first 50,000 lines of the log file from the Reporting Server browser interface. If you wish to view lines beyond 50,000, you can open or save the file in a text editor.

Controlling the HLIPRINT Transaction Log

You can control the size of the HLIPRINT transaction log by using the `hliprint_max_lines` parameter in the configuration file, `edaserve.cfg`. The parameter defines the maximum number of lines to be included in `hliprint.log` before it is archived and a new `hliprint` log is started.

Each time the line limit is reached, the current HLIPRINT log file is archived to `HLIPRI01` and the new file starts with an additional identification line. This line indicates that the file was truncated and looks like the following:

```
continued from hlipri01
```

Archiving also takes place the first time the FDS listener is activated and the HLIPRINT log from the previous run is found in the server directory.

Procedure: How to Control the Size of the HLIPRINT Log From the Reporting Server browser interface

1. Go to the *Workspace* page.
2. Expand the *Special Service and Listeners* folder in the resources pane.
3. Right-click *FDS*, and click *Properties*.

The Special Services Configuration page opens.

- Expand the *Advanced* section by clicking the arrow.

The screenshot shows a window titled "Special Services Configuration". It has two main sections: "Basic" and "Advanced".

- Basic section:** Contains three fields:
 - NODE:** A text field containing "FOCSU".
 - PORT:** A text field containing "8118" with "(required)" to its right.
 - HOST:** An empty text field.
- Advanced section:** Contains one field:
 - hliprint_max_lines:** A text field containing "1000".

At the bottom of the window are two buttons: "Save and Restart Special Service" and "Cancel".

- Enter a value for:

`hliprint_max_lines`

Is the maximum number of lines to include in the `hliprint.log` before it is archived and a new `hliprint` log is started. You can limit the size of the log file by setting the maximum number of lines to a value other than zero. (0 means an unlimited number of lines can be included in the file.)

There is no minimum, but the recommended value is at least 1000 lines. If the limit is too low, the FDS listener may experience performance problems due to excessive archiving.

- To register this setting, click *Save and Restart Special Service*.

Controlling HLIPRINT Output

The Server Administrator can control `hliprint.log` output by manually adding the `EDAHLIP` variable, with a value, to the `edaserv.cfg` file, in the form `edahlip=value` where valid values are:

`none`

No `hliprint.log` is created. Turning logging off will increase performance, however, this is not recommended, because it makes it harder to diagnose problems.

echo

The hllprint.log is created and contains basic transaction information (one line per transaction; line length 80 characters).

stat

The hllprint.log is created and contains basic transaction plus timing information (one line per transaction; line length is 133-characters).

The default value is stat.

Tracing TIBCO WebFOCUS Reporting Server Activity

You can activate traces at server startup or while the server is running. It is recommended to start the server with traces active so that a full set of traces can be sent to Customer Support Services to help diagnose issues that arise. If you need to contact Customer Support Services for assistance, you can use a facility called savediag to expedite the transmission of server information. For details about using savediag, see [Gathering Diagnostic Information for Customer Support Services](#) on page 794. Refer to the *TIBCO WebFOCUS® Installation and Configuration* manual for your platform for more information on starting the server with traces active.

To activate traces for a running server, a server administrator can quickly enable or disable traces at any point in a work process by clicking, *Enable (or Disable) Traces* from the User quick link menu on the Reporting Server browser interface menu bar.

The server administrator can use the Reporting Server browser interface to initiate and control a wide range of traces, and dynamically turn traces on or off for a running server. To see a list of traces, go to the *Workspace* page. Open *Logs and Traces* on the resources pane, right-click *Traces*, and click *View*.

The Reporting Server browser interface displays the following types of traces: Workspace Manager, Data Services, Java Services, Special Services, and Listeners. (The traces that are active vary depending on the types of requests that have been made against the server.)

Right-clicking a trace displays the following options:

- View.** Displays the trace in the right pane.
- Download.** Enables you to open the trace in a local editor or to save it to a local disk.
- Delete.** Enables you to delete the selected trace file. On Windows, a trace file cannot be deleted if it is still in use by a process.

Note that the selected trace file is displayed with any active filtering applied.

Once traces have been configured and turned on by a server administrator, they can be viewed by users with the General privilege SRVLG (View Server Logs and Traces, Create Savediag), which can be configured from the Access Control page.

Procedure: How to Access the Last Agent Trace

To access the last agent trace:

1. Go to the *Workspace* page.
2. Click *Diagnostics*, point to *Download*, and click *Last Agent Trace*.

The last agent trace opens.

Reference: SET TRACEUSER Commands

If you want to turn dynamic traces on or off for one user, you can include the command SET TRACEUSER={ON|OFF} in a user profile or in an individual procedure.

You can also use the SET TRACEUSER=*tracename* command turn dynamic traces on and send the trace to a named file. The file must specify a .trc extension and use a full path name for server traces so the traces are retained after server execution.

Configuring Traces

You can configure traces to limit their size and select the components to activate.

1. Go to the *Workspace* page.
2. Open *Logs and Traces* on the resources pane, right-click *Traces*, and click *Configure*.

The Traces Configuration page opens.

The screenshot shows a window titled "Traces Configuration" with the following content:

- Setting will be activated for next workspace run
- Set number of lines to limit the trace size:
- (leave blank for no limit)
- Set column width of trace:
- (Default is 230)
- Choose the trace components to activate:
- Buttons:

3. Optionally, set the trace size limit to a specific number of lines. After the limit is reached, trace data is dumped to an alternate file with the extension TRB. Logging continues into the trace file with extension TRC. For related information, see [Limiting the Size of a Trace File](#) on page 771.

A blank entry indicates no trace size limit.

4. Optionally, change the column width of the trace file to reduce line wrapping and make the trace file more readable. The default column width is 230 characters.
5. Select an option from the *Choose the trace components to activate* drop-down menu. The selections you make here are written into a file called `ibitrace.fex`. The settings take effect on the next server start/restart.

Note that Customer Support Services may request that you change from the default trace setting to a custom setting to facilitate their research.

Default Components

The Default trace components include all trace levels, with the exception of Reporting and Miscellaneous Services low level traces.

All Components

- All traceable server components.

Typical Components

- CEH
- NWH2
- PRH
- SQLAGGR
- STMTRACE

Custom Components

This option allows you to select the specific components to activate.

The components are categorized as follows (each category has a list of trace options within it):

- COMMAND PROCESSING
- COMMUNICATIONS and SYSTEMS
- COMMON ADAPTER LAYERS AND SERVICES

- SQL TRANSLATOR
- SQL DATA ADAPTERS
- OLAP DATA ADAPTERS
- NON-SQL DATA ADAPTERS
- LEGACY DATA ADAPTERS
- REPORT SERVICES
- NLS SERVICES
- ODIN INFORMATION
- SMART MODE
- MAINTAIN PROGRAMMER TRACES
- MAINTAIN INTERNAL SERVICES
- MISCELLANEOUS SERVICES
- TRANSACTION COORDINATOR
- XFOCUS
- USE-WHERE

6. Click the *Save* or *Save and Restart Server*.

Tip: If you wish to restore default settings or choose typical settings, select these options, then click *Save Settings*. The selected group will be checked when you reopen the Custom Components list.

Reference: **Limiting the Size of a Trace File**

You can use set the size limit of the trace files. The trace size limit specifies the maximum number of lines in the trace files. The value is written into the IBITRACE.FEX file as SET TRACESIZE = *value* and requires a server restart to activate. After the size limit is reached for a particular trace, the trace data is redirected to an alternate file with an extension of .TRB. When the size limit is reached in the .TRB version of the file, then writing resumes with the .TRC version of the file.

A trace cycles between the two matched sets of files to alleviate maxing out a disk drive with large individual traces. Each cycle starts a fresh file (versus an append). If there are intermediate traces that need to be saved, it is up to the site to monitor and implement specific methods to save traces based on their needs and capabilities. For example, copying to an archive disk or secondary machine and possibly using compression.

The SET TRACESIZE=# setting must be used in conjunction with a SET TRACEON= value. If the IBITRACE.FEX file is changed or created using an editor, ensure that there is also a SET TRACEON=value setting. The existence of the IBITRACE.FEX file only sets what is traced when tracing is active, but does not control if tracing itself is active.

Reference: Displaying Microseconds in Trace Timestamps

Displaying microseconds may be useful if trying to match very precise event times between traces or study the performance of specific events when multiple events occur within a given second on a normal trace.

The setting syntax is SET TRACESTAMP=TLEFT and may be used in a focexec to activate in a specific trace or in the IBITRACE.FEX file to activate at server start.

There is no Reporting Server browser interface feature on the Trace Configure page to activate TRACESTAMP in IBITRACE.FEX, but the line can be added by opening the Configuration Files folder in the Reporting Server browser interface and using the edit against the Trace Profile object.

The SET TRACESTAMP=TLEFT setting must be used in conjunction with a SET TRACEON= value. Always ensure that there is also a SET TRACEON=value setting in the IBITRACE.FEX file when editing. The existence of the IBITRACE.FEX file only sets what is traced when tracing is active, but does not control if tracing itself is active. The SET TRACESTAMP=TLEFT feature also negates the ability to properly use Reporting Server browser interface Trace Filtering for specific events, do not use filtering with this feature active.

Filtering Traces

Filters allow you to limit the types of errors or components that are included in the traces.

To filter traces:

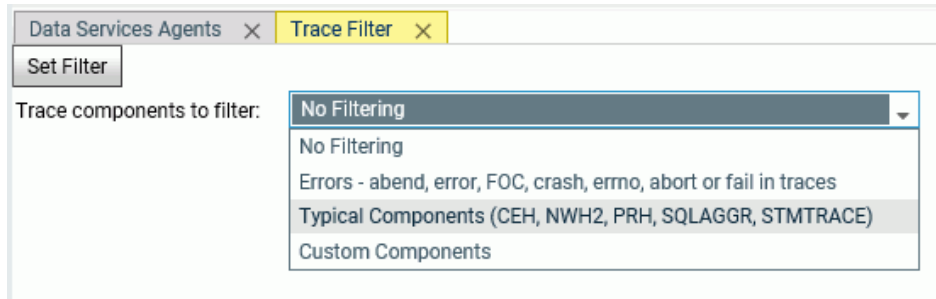
1. Go to the *Workspace* page.
2. Expand *Logs and Traces* on the resources pane, right-click *Traces*, and click *Filter*.

The Traces Filter page opens.

3. Select an option from the *Trace components to filter* drop-down menu:

No Filtering (the default) displays all trace information.

- Errors* displays only errors, crashes, errno, abort, or fail conditions.
- Typical Components* displays trace information from typical components.
- Custom Components* displays trace information from components which you select.



4. Click *Set Filter*. The filters take affect when you enable and run traces.

Procedure: How to Capture Traces in a Script File

Capturing traces in a script file enables you to isolate server problems. Specifically, you can enable this type of custom trace, then reproduce the problem and capture the input commands issued from a front-end tool to the server. Customer Support Services can then isolate debugging for the server only, instead of having to recreate a whole complex environment.

1. Go to the *Workspace* page.
2. Expand *Logs and Traces* on the resources pane, right-click *Traces*, and click *Configure*.
3. Select *Custom Components* from the *Trace components to filter* drop-down menu.
4. Uncheck the box on the menu bar to turn off all selections.
5. Choose *AJ - NWHSIM/1* under *Communications and Systems*.
6. Click *Set Filter*.

The trace setting will be active the next time the server is started.

Showing the Session Log

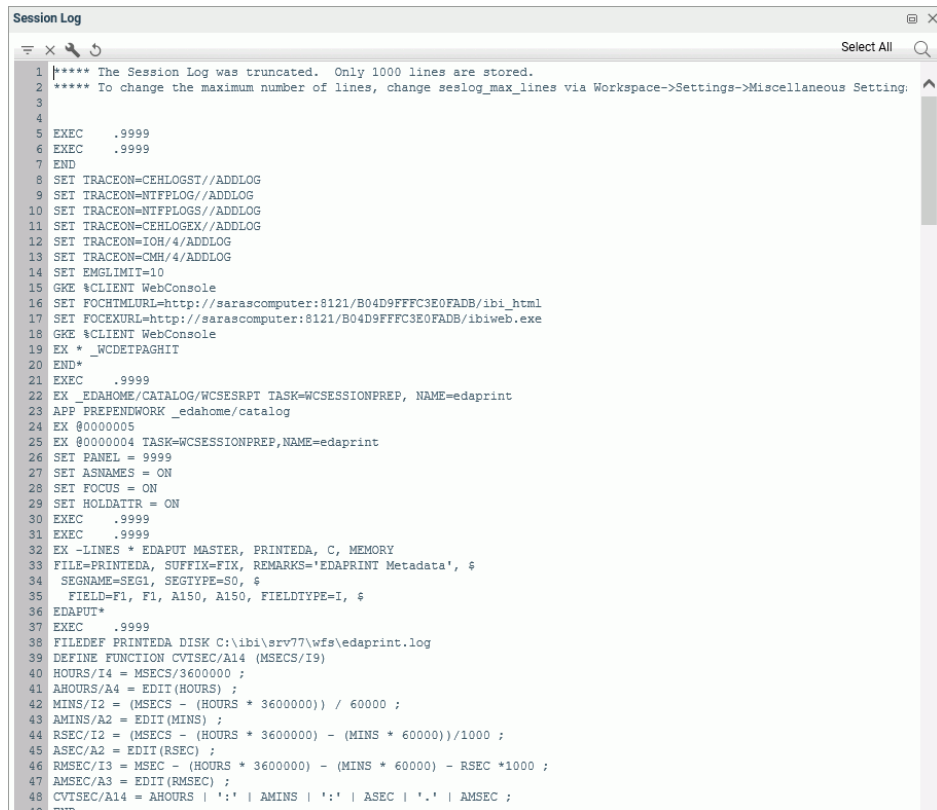
The Session Log records the activity that takes place between the current Reporting Server browser interface session and the server. For example, when a procedure or Data Migrator job is run, the Session Log records the WebFOCUS code and SQL statements that are sent from the console to the server for execution on the server.

The log also records the response that the server returns, in addition to the answer set, such as error messages and informational messages.

To show the Session Log:

1. On the Reporting Server browser interface *User* menu, click *Session Log*.


The Session Log opens, as shown in the following image.



2. Click the *Refresh* button on the left of the Session Log menu bar to display the current contents of the log.
3. Optionally, click the *Filter* icon to select only those commands or data that you want to include in the Session Log. You can select:
 - Focus, to display WebFOCUS commands sent to the server for execution.
 - Sql/mdx, to display SQL or MDX code that is generated by an adapter and sent to a relational database or OLAP database, respectively.
 - Transform, to display the result of transforming the user-supplied code in a procedure, into WebFOCUS code that is supported on the server. This option is useful for Data Migrator jobs that require transformation.

- Errors, to display error messages returned from the server.
- Messages, to display informational messages returned from the server, for example, NUMBER OF RECORDS IN TABLE, SVG FILE SAVED, and others.
- Timestamp, to display the date and time at which a command or message was exchanged. For example, if you select *Errors* and *Timestamp*, the Session Log displays the time at which an error message was returned from the server.
- No Delay, to display the Session Log while a request is still running, not after it has finished running. This option is recommended for long-running requests.

In the following image, the check boxes for Focus, Messages, and Timestamp are selected to include the display of WebFOCUS commands and messages, and to include the date and time of each entry in the log.



When you filter the Session Log, its contents automatically refresh to reflect the last criterion that you specified.

4. Optionally, you can do one of the following:
 - Click the *Clear Messages (X)* button on the left of the Session Log menu bar to remove the current contents of the log.
 - Click the *Refresh* button on the left of the Session Log menu bar to update the log.
 - Use standard editing functions, such as:
 - Ctrl+A, Ctrl+C, and Ctrl+V to select all the contents of the log, copy it, and paste it into another file.
 - Ctrl-click to highlight a line.
 - Shift at the beginning of a block of lines, and Shift-click at the end of the block of lines, to highlight the lines.
 - Click the *Options* menu (down arrow) on the right side of the Session Log menu bar and click *Text Editor Options ...* to open the Editor Options dialog box. The options *Enable Code Folding*, *Highlight Active Line*, and *Show Line Numbers* are checked by default.
 - Click the *Select All* icon (square surrounded by dots) on the right side of the Session Log menu bar to select all lines in the Session Log.
 - Click the *Find* icon (magnifying glass) on the right side of the Session Log menu bar to open a search text box.
5. When you have finished reviewing the messages in the Session Log and want to close it, click *Close (X)* on the top right of the Session Log window.

Changing the Size of the Session Log Buffer

By default, the Session Log holds 1000 lines. You can increase or decrease the number of lines.

To change the size of the Session Log buffer:

1. Go to the *Workspace* page.
2. Click *Settings*, point to *Workspace SET*, then click *Miscellaneous Settings*.
3. On the *Miscellaneous Settings* page, locate the *seslog_max_lines* field, and type the number of lines to assign to the buffer for the Session Log.
4. Click *Save*. You do not need to restart the server for the new setting to take effect.

Procedure: How to Show the Log for a Web Session

You can display a log of the activity that takes place in an individual Web Session.

1. Go to the *Workspace* page.
2. Click *Monitor*, then click *Web Sessions*, or expand *Special Services and Listeners* on the resources pane, right-click *TCP/HTTP*, and click *Web Sessions*.

The Performance Report for all Web Sessions is displayed.

3. On the report, right-click a row for an individual Web Session, and from the menu, click *View Session Log*.

The log for the selected Web Session opens.

z/OS-Specific Troubleshooting for the Unified TIBCO WebFOCUS Reporting Server

If you are working in the z/OS environment, you should be aware of several diagnostic techniques that are described in this topic.

- Server Traces and Saved Diagnostics
- U4039 Abends
- System Dumps
- Freeing Data Sets Allocated to the Server

Traces and Saved Diagnostics

To generate a server trace:

1. Turn tracing on from the Reporting Server browser interface or by issuing the MODIFY operator command:

```
F <iway_server_jobname/started task>, -tracoon
```

2. Alternatively, on HFS you can run the ITRCON JCL member.
3. Reproduce the problem.
4. Submit the ISAVEDIA member to produce the diagnostic information.

For PDS deployment, diagnostic information is generated at the JES log output of the server.

For HFS deployment, a directory called *sdnnnnn* is created under your configuration directory (for example, */ibi/server53/ffs/sd123456*). Diagnostic information is placed in this directory. Verify that you have access to this directory.

Do not change anything in the EDAENV DD statement, as changes could prevent the correct information from being copied to your directory.

U4039 Abends

In some instances, the server may abend with a U4039 code. This is a generic abend.

Solution: Determine what caused the abend by checking the

file, SYSOUT DDname and the MVS system log. If you suspect this is a bug, follow the instructions in [z/OS-Specific Troubleshooting for the Unified TIBCO WebFOCUS Reporting Server](#) on page 777 and [System Dumps](#) on page 778 in order to collect information to send to Customer Support Services.

System Dumps

To generate a system dump:

1. Allocate DDNAME SYSMDUMP pointing to the data set with the following DCB parameters:

```
RECFM=FB,LRECL=4160,BLKSIZE=4160
```

2. To get the first dump, add the parameter FREE=CLOSE to your DD statement. The DD statement should appear as follows:

```
//SYSMDUMP DD DISP=SHR,DSN=MYID.EDAPTH.SYSMDUMP,FREE=CLOSE
```

3. To get the last dump, the statement should appear as follows:

```
//SYSMDUMP DD DISP=SHR,DSN=MYID.EDAPTH.SYSMDUMP
```

Only two IDs must have privileges to write into this data set: ISERVER and IADMIN. General server users DO NOT need read or write access to the SYSMDUMP data set.

4. To prevent abendaid from intercepting the dump, add:

```
//ABNLIGNR DD DUMMY
```

5. To prevent Language Environment from intercepting the dump, specify:

```
EDADUMPOPT=UAIMM in EDAENV DD
```

This enables you to get more accurate information reflecting the moment the abend actually occurs.

6. Save the entire job output for the server (including JES logs).

Freeing Data Sets Allocated to the Unified TIBCO WebFOCUS Reporting Server

A z/OS operator can issue modify commands to free DDNAMES (or dsnames) allocated to the Server. Both global allocations (made at the server ISTART JCL) and local ones (DYNAM ALLOC commands issued by user tasks) can be freed. This procedure is useful if you need to free an allocation to run a batch utility overnight, without restarting the server.

How to Free a Data Set From the MVS System Console

To free a single DDNAME:

```
F <iway_server_jobname/started task>,DYNAM FREE FI <ddname>
```

To free a single dsname (all occurrences in the server):

```
F <iway_server_jobname/started task>,DYNAM FREE DA <dsname>
```

To free multiple DDNAMEs, passing a pattern (free all ddnames starting with AB):

```
F <iway_server_jobname/started task>,DYNAM FREE FI AB*
```

To free a multiple dsname (all occurrences in the server), passing a pattern (free all allocations of data sets starting with IWAY.VSAM*):

```
F <iway_server_jobname/started task>,DYNAM FREE DA IWAY.VSAM*
```

A message will be issued in the iway_server JESMSG LG indicating if the command was processed successfully or not.

success:

```
@DYNAM COMMAND SUCCESSFULLY PROCESSED Rc=0
```

failure:

```
@DYNAM ERROR: IKJ56225I DATA SET IWAY.TEST ALREADY IN USE, TRY LATER
```

Once the data set is freed and the batch utility runs, the operator can reallocate the data set to the server, using the same modify command, but this time issuing DYNAM ALLOC, instead of DYNAM FREE.

Example: Freeing the Allocated Data Set

Suppose ISTART JCL (jobname IWAYP) has the following allocation:

```
//VSAMFILE DD DISP=SHR,DSN=VSAM.FILEA.CLUSTER
```

The operator can free this file using the command (from MVS console):

```
F IWAYP,DYNAM FREE VSAMFILE
```

Later, after the batch utility runs, the operator can re-allocate the file to the server by issuing the command:

```
F IWAYP,DYNAM ALLOC F VSAMFILE DA VSAM.FILEA.CLUSTER SHR REU
```

All valid DYNAM ALLOC and DYNAM FREE syntaxes are supported. For more information on DYNAM command, please refer to the appropriate WebFOCUS manual.

z/OS-Specific Troubleshooting for z/OS HFS Deployment

The troubleshooting techniques described in this section are available for a z/OS HFS deployment.

- Recovering Orphaned Shared Memory Segments
- INSUFFICIENT AUTHORITY TO GETPSENT Messages

Recovering Orphaned Shared Memory Segments on z/OS

When the server is restarted after an abend, an orphaned shared memory segment is often found. A server start will try to clear the memory segment so that it may continue the start process. The accompanying error message (which also displays even if the server *self-recovers*) is:

```
Shared memory remains but WSM process is gone
```

Self-recovery is often successful and the message may be ignored. However, if the server still fails to actually start, you can use manual commands to find and remove the shared memory. The associated directory location that is used as the shared memory key will not actually be available until a machine IPL. To avoid waiting for a machine IPL, you can set EDASHARE manually in ISSTART EDAENV DD and the ISTOP JCL to point at a new directory location.

1. To display the orphaned memory objects, issue the command:

```
IPCS M -X
```

2. To remove the orphaned memory objects, issue the command:

```
IPCR M
```

IPL will eliminate the orphaned shared memory segment.

3. If you cannot IPL the machine, temporarily add the following EDASHARE setting (with the client directory path) to the ISTART EDAENV DD.


```
000050 //EDAENV      DD      *
000051 TZ=EST5EDT
000053 EDASHARE=/ibi/srv90/ffs/web
```

where:

web

Is a newly created directory. It can be any name.

This will point the server to a new directory to be used as the shared memory segment key.

This syntax should also be added to the ISTOP JCL, but should be removed from both JCLs after IPL.

INSUFFICIENT AUTHORITY TO GETPSENT Messages on z/OS

INSUFFICIENT AUTHORITY TO GETPSENT messages may appear in JESLOG. For more information about what causes these messages, see the IBM APAR II11813.

The APAR recommends issuing one of the following RACF commands to work around this problem:

```
SETROPTS LOGOPTIONS (NEVER(PROCACT))
```

```
SETROPTS LOGOPTIONS (DEFAULT(PROCACT))
```

However, when a non-superuser in the OMVS shell issues the command `ps -ef`, the following security message is repeated in SYSLOG:

```
ICH408I USER(default) GROUP(dgltgrp) NAME(bpxdefaultuser) 060
CL(PROCACT ) INSUFFICIENT AUTHORITY TO GETPSENT
```

This does not indicate an error. It is an informational message issued because of RACF LOGOPTIONS settings. The `ps -ef` command is a request to show all processes that the requester is authorized to see, but a non-superuser is allowed to see only his or her own processes.

Recording and Reproducing User Actions

You can use the Reporting Server browser interface Record and Playback tools to record an exact sequence of user actions applied through a browser, and then reproduce it (play it back) to evaluate server response times in an environment that simulates single or multiple users under the same or different conditions. Files generated by record and playback, known as HTI scripts (HTTP internal scripts), are stored in the scale directory of the server (parallel to APPROOT). The file name can be edited before recording starts. These scripts, along with server traces, are used for testing and analysis of problems.

Record/Playback is an iterative process in which you can repeat your tests under varying conditions until you achieve the throughput and response time you require.

Tip: It is often useful to employ two browser windows for stress testing using the Record/Playback tool, one to turn the recording on and off and enter playback parameters, the second to capture the actions you wish to record. (Note that if you do not use two sessions, you will have to edit the request being stressed to remove the record/setup actions.)

In order to be able to capture these actions, you must properly configure your LAN setting.

Here is an overview of the steps:

1. Set your browser to use the Reporting Server as a proxy server. This will allow you to capture user actions in a web browser. For details, see [How to Configure Your Browser to Use the Reporting Server as a Proxy Server](#) on page 782.
2. Start the script recording.
3. Run the request you want to test.
4. Stop the recording. You will be able to reuse this script with varying test conditions.
5. Play back the script you recorded, setting stress parameters, such as the number of threads (simulated users) and the intervals at which the request will be run.
6. Evaluate server performance using the Playback log, the output (.hto) files generated by the test, and stress test performance reports.
7. Repeat steps 2-7 as many times as necessary to simulate and evaluate the conditions you want to test.

Note: The Report/Playback feature is not currently supported on the PDS Deployment Server. Therefore, options are disabled on the Reporting Server browser interface.

Procedure: How to Configure Your Browser to Use the Reporting Server as a Proxy Server

When using a Reporting Server as a proxy, you can record user actions that are directed to any web site. In order to be able to record the sequence of actions that a user will complete using the browser, you must first enter the proper LAN settings.

1. From the Internet Explorer toolbar, select *Tools* and then *Internet Options* from the browser menu bar.
2. Select the *Connections* tab, then click the *LAN Settings* button. The Local Area Network (LAN) Setting window opens.

Tip: Before you make the required changes, note your standard settings. You will want to revert to them later.

3. Click the *Use a proxy server for your LAN* check box.

4. Enter the host name in the Address box. (This identifies the machine on which the client software is installed.)
5. Enter the port number for which the client is configured.
6. Click *OK* to return to the Internet Options window.
7. Click *OK* again to complete the task.

Leave the browser open since you will use it to display the Reporting Server browser interface.

Procedure: How to Stress Test Server Performance

The simplest way to record an HTI script is by running a stress test on a procedure. For information, see [How to Run a Stress Test](#) on page 566.

1. Go to the *Scalability* page.
2. Click *Record*.

The Recording of HTI Script File pane opens.

3. Specify a script name and the directory in which you want to store it.

Recording is Off

Proxy Mode:	Off
Location:	C:\ibi\scale
Script Directory:	test0001 (type in a new one or select one from the list)
Script Name:	rec0002
<input type="checkbox"/> Append to existing file	
<input type="checkbox"/> Do not record requests to get css, js, gif, bmp, jpg, class, jar, or png files	
<input type="button" value="Start"/>	

A default directory name appears in the Script Directory drop-down list. You can use the default name, edit it, or choose an existing directory from the list.

4. Enter a script name or use the default name. This is the name of the file that will contain the script.
5. By default, a new HTI script is generated when you start a recording session. Click the *Append to Existing* check box if you prefer to append a recording to the existing HTI script.

6. By default, the recording captures all user actions. Click the *Do not record requests to get css, js, gif, bmp, jpg, class, jar, or png files* if you wish to exclude requests for these operations from the recording.
7. Before you start the recording, you must get into position to execute the set of actions you want to test. Open a second browser and ensure that the correct LAN settings are in effect. (For details, see [How to Configure Your Browser to Use the Reporting Server as a Proxy Server](#) on page 782.)
8. Return to the first browser window, where the Recording of HTI Script pane is still open, and click *Start* to begin recording.
9. Return to the second browser window and perform the actions you wish to record.
10. When you have captured all of the required actions, return to the first browser window, where the first the Recording of HTI Script pane is still open. Click *Stop* to end the recording.
11. Right-click the application folder with the script, and click *Refresh* to display the new script.
12. Since you have completed the recording task in the second browser window, it is a good idea to remove the LAN setting you entered expressly for this purpose. (For more information, see [How to Configure Your Browser to Use the Reporting Server as a Proxy Server](#) on page 782.) You can now close this window.
13. You are now ready to perform your stress tests. On the resources tree, expand the *Individual Tests* folder, expand the script directory you named in step 3, and right-click the script name you entered in step 4.
14. Click *Start* from the menu. The Playback of HTI Script Files pane lists the available test parameters.

You can choose basic and advanced test options from the Playback of HTI Script Files pane. For details about these options, see [Parameters for Playback of HTI Script Files](#) on page 787.

The options on the ribbon are Recording, Script, Stop all Tests, and Refresh.

For example, you might wish to perform a stress test twice with the following conditions:

Test 1: Use the default *Number of Threads* (1) to simulate one user and the default *Interval Parameters* (.5) second to indicate how often you want the request to run. Enter a *Test Description*, such as *Baseline with one user*, to be able to distinguish among the tests. Click the *Advanced* tab and accept the default format of *one file per thread*. (In this case, you are only testing one thread.)

Click *Start* to run the test.

Test 2: Set *Number of Threads* to 30 to see how performance would be affected if 30 users ran the request simultaneously (with a given number of agents that may or may not be less than 30). This setting is key to stress testing. Adjusting the number of threads allows you to assess how many requests the server will be able to handle at a given point in time. Increasing the number of threads puts more load on the server(s). In practice, users may want to start out with a smaller or larger number depending on their installations. Keep the *Interval Parameters* as .5 seconds. Enter *Sample 30 Users* as the Test description.

Expand the *Advanced* tab. This time select *one file per request* (rather than per thread) in the Results field and set the *Keep alive* parameter to 30 to indicate how long, in seconds, you want the script to run. This option enables enough statistics to be generated to make a comparison among server response times. (You may need to try different settings to find the optimal length of time to run the script.)

Click *Start* to run the test.

In each case, two files are stored in the *ibi/scale/performance* directory:

- A summary log is displayed in a separate window. The log provides information about the number of threads that ran (indicated by *thread=*), the number of requests serviced (indicated by *request=*), how long it took to service each request (indicated by *resp=*), and the time that the processing began (indicated by *start=*). It also reiterates your input and summarizes the test processing.
 - An *.hto* output file is displayed in the Playback, Individual Tests folder, under the specified script directory and script name. This file indicates whether the client/server communication worked properly.
15. You can perform the following additional functions from the Individual Tests folder. Right-click a script name and choose:
- Show requests.** Lists all submitted hti requests.
 - Show result.** Shows the generated output.
 - Edit script.** Opens the hti script file in a text editor. You can save the script or refresh from the disk.
 - Start.** Opens the Playback of HTI Script Files pane, where you can revise parameters and rerun the script.
 - View Report.** Produces the same output as the log file in report format.
 - View Log file.** Reopens the log file.

- View Graph.** Produces the same output as the log file in graph format. The graph is particularly useful for comparing the average response time of the server with your specified Maximum Response time to determine if you have met your performance goals.
 - Delete Result.** Deletes the generated output but retains the script for reuse.
 - Delete Script and Result.** Deletes the generated output but retains the script for reuse.
16. Once the script is run, you may also want to open the Performance Reports folder and choose:
- Basic Report to see basic stress test statistics.
 - Extended Report to see an expanded group of stress test statistics.
 - Graph Report to a bar chart of stress test statistics.
- The right-click options for the reports are View, Save in EXCEL, and Open.
17. Typically, you would run additional tests with different input parameters in order to compare statistics under different simulated loads conditions. Based on these comparisons, you could determine the best server scenario for your needs.

Procedure: How to Refresh or Remove Script Directories

You can expand the Individual Tests folder and click a specific directory. A pop-up menu enables you to remove the selected directory or refresh it without having to refresh all of the other test directories.

1. Go to the *Scalability* page.
2. From the resources pane, expand the *Individual Tests* folder (if not already expanded).
3. Right-click a directory and choose either *Refresh* or *Delete Script Directory* from the menu.

Procedure: How to Display Percentile Statistics in Results

A percentile is the point in a distribution of scores below which a given percentage of scores is found. When you perform a stress test you can calculate percentile statistics for processing time by specifying the percentage against which the percentile value will be derived.

You request percentile statistics from the Playback HTI Script Files pane, which you can access in either of the following ways:

As a diagnostics option:

1. Go to the *Scalability* page.
2. On the resources pane, expand a directory folder, right-click a script name, and click *Start* from the menu. The Playback of HTI Script files pane lists Basic and Advanced test parameters. Enter basic test conditions: *Number of Threads* and *Interval Parameter*.
3. Expand the *Advanced* section. A variety of additional test options are displayed.
4. In the *Show percentile statistics* field, enter a number N , between 1 and 99, to calculate the processing time in your stress test distribution against which $N\%$ of the other processing times are lower. For example, if the distribution of 6 processing times in a stress test includes the following values: 1-1-3-3-4-5, and the percentage value for *Show Percentile Statistics* is set to 50 (%), then the 50th percentile processing time will be 3 because that is the lowest value in the distribution against which 50% of the processing times are lower.
5. Click *Start* to run the test with the specified condition(s).

The Percentile result, along with other test results, appears in a separate window.

When stress testing a particular procedure:

1. Go to the *Applications* page.
2. Expand an application directory, right-click a procedure, and select *Run Advanced*, then *Run Stress*. The Playback of HTI Script Files opens. This pane lists the available test parameters.
3. Enter basic test conditions: *Number of Threads* and *Interval Parameter*.
4. Expand the *Advanced* section, and proceed as described in steps 4 and 5 above.

Reference: Parameters for Playback of HTI Script Files

Following is a list of general playback parameters that are available.

Basic:**Number of Threads**

A positive number of threads to be used when playing the script. Each thread represents a single client (user).

Interval Parameter

Specifies timing in a multiple-client (multi-threaded) playback. Possible formats of this value are m or m, n, k , where m is the number of seconds between each client startup, and after each n started clients, an interval of k seconds is used instead.

Test Description

Enables you to create notes at playback time that will be displayed in the Extended Report.

If you reset statistics, all information about previous runs is deleted, including any Test descriptions.

Tip: This description provides a useful way to distinguish results as you test different combinations of options.

Advanced:

Compare option

Determines whether the response data received by all clients (threads) are identical. It uses a binary comparison mode and writes the comparison results to a playback log file.

Immediate processing

Ignores every SLEEP and WAIT statement in the HTI script. Issues HTTP requests without any delay between them.

Show percentile statistics

A percentile is the point in a distribution of scores below which a given percentage of scores is found. When you perform a stress test you can calculate percentile statistics for processing time by specifying the percentage against which the percentile value will be derived.

Enter a number N , between 1 and 99, to calculate the processing time in your stress test distribution against which $N\%$ of the other processing times are lower. For example, if the distribution of 6 processing times in a stress test includes the following values: 1-1-3-3-4-5, and the percentage value for Show Percentile Statistics is set to 50 (%), then the 50th percentile processing time will be 3 because that is the lowest value in the distribution against which 50% of the processing times are lower.

For related information, see [How to Display Percentile Statistics in Results](#) on page 786.

Size statistics

Displays general output file size statistics for multiple-client playback in the log.

These are useful for analyzing Playback results with a large number of threads. Size should be consistent between similar runs. A size change would generally indicate an execution failure.

Simulate browser caching

Enables browser-style simulated caching of static pages (currently, css, js, gif, bmp, jpg, class and jar files). Each thread (simulated user) maintains its own cache of requests.

When a cached static URL is requested again by the same user, the Stress tool skips the request instead of downloading the URL again.

Submit as deferred

Executes a request using deferred, rather than immediate, logic.

Save playback statistics

Choose this option to make statistics for the current run available on reports.

The statistics are saved internally in the \$EDACONF/catalog directory as rec0001.hti, rec0001.log, and rec0001_0001.hto.

For example, if the script generated in a recording session has the name rec0001.hti, then the run log and statistics are stored in the rec0001.log file and the output generated by each of three threads is rec0001_0001.hto, rec0001_0002.hto, rec0001_0003.hto.

Traces

Sets the Playback utility traces on or off. A trace file (TRC) is then created for each thread.

Keep alive for

Causes every thread to repeat a request for a specified time, in seconds. This option enables the server to generate enough statistics to make a comparison among response times.

Result

The HTTPST module that records and plays back Reporting Server browser interface sessions can simulate multiple client access to the Reporting Server browser interface on playback (that is, a multi-threaded application).

Each thread does what is recorded in the .hti file. The server reply is recorded in the .hto file. However, you can control the output provided by selecting one of the following options:

- One file per thread (default).** Displays output in .hto files. The server generates a summary .hto file that reports on the success or failure of the server/client communication in the format: For example:

```
Number of successfully run threads :5
```

You can access this file by selecting *Playback*, *Individual Tests*, and then the *script_directory* folder in the format

```
scriptname--nnnn
```

where:

nnnn

Is a number that is appended to the file name. It indicates the size of the generated file.

This option also generates a separate .hto file for each thread being tested; these files contain the code that renders the report. You can access these files by selecting *Playback, Individual Tests*, and then the *script_directory* folder in the format

scriptname_thread#_filesize

and view the contents in a text editor.

- ❑ **One file per request.** Displays a single .hto file in the Playback>Individual Tests folder. This is the summary file that reports on the success or failure of the server/client communication. For example:

Number of successfully run threads :5

You can access this file by selecting *Playback, Individual Tests*, and then the *script_directory* folder in the format

scriptname--nnnn

where:

nnnn

Is a number that is appended to the file name. It indicates the size of the generated file.

In addition to the summary file, the output for each request is written to the *ibi/scale/performance* directory in the format

scriptname_thread#_request#.ext

where:

ext

Is determined by the nature of the request (for example, HTML), and by its success or failure (failure is often represented as a .txt file). For example:

rec0001_0002_0041.html

where:

`rec0001`

Is the script name.

`0002`

Is the thread number.

`0041`

Is the request number.

Additional requests in the thread would be represented as:

```
rec0001_0002_0042.html
rec0001_0002_0043.html
```

Server

Sends request to a server defined in the `odin.cfg` file.

Use HTTP proxy

Send all requests through an HTTP proxy server. For example, `www.myproxy.com:3192`.

Save output to another directory

Saves output files, including traces, to a specified location (for example, `c:\myfiles`).

Reference: Additional Recording Hints

- When recording is on, the HTTP listener records every CGI request. It is important to ensure that no one else is submitting requests to the same listener.
- Do not try to record a script from the middle of a persistent session. Always try using user ID and password, then recording the first request.

Reference: Playback Files

Playback log files are generated during script playback. Each file contains general information about playback processing, such as thread startup and termination, connection errors, and comparison results.

- Log file names are composed by appending `.log` to the base file name of the HTI script used to generate the log. Two additional log files, `stdout` and `stdlog`, may be generated to report critical playback failures.
- Playback HTO files are generated when the script playback is complete. Every file contains data received by a thread.

HTO file names are composed of the script file name, followed by the thread number and the .hto extension.

Troubleshooting the Reporting Server Browser Interface

If the web browser has problems accessing the Reporting Server browser interface, you may see the following messages.

Error Message	Troubleshooting Tips
Internet Explorer cannot open the Internet site http://address:http_service. A connection with the server could not be established.	Ensure that the Workspace Manager is running or else contact the local system administrator.
There was no response. The server could be down or is not responding. If you are unable to connect again later, contact the administrator of the server.	Ensure that the Workspace Manager is running or else contact the local system administrator.

TIBCO WebFOCUS Reporting Server Processes

The server uses a series of processes to accomplish various tasks. Some processes have specialized administrative functions, such as workspace management and logging, while others do end-user tasks. System administrators may want to know what processes the server creates and how many of each. The administrator may want this information for machine sizing or simply to know what processes to consider or ignore when a problem occurs.

Note that each operating environment uses operating system-specific commands to view processes. The processes are described generically, not how they would appear in a system listing.

Server Work Space Manager Daemon

The server workspace is control by a daemon process known as PTH. There is one process per server.

Deferred Requests Special Services Daemon

The server has several types of special services. The DFM daemon is used for deferred (DFM) and Scheduler requests. There is one process per server.

Server Log Daemon

The server log (edaprint.log) is now written using a single daemon process. Previous versions wrote to the log independently from the various processes that might write to the log. In very rare and specific cases, the prior methods were problematic, hence, the change to the use of a daemon. As a result, starting a server brings up this additional server log daemon (edaplog).

Server Check Up Daemon

The server has always had a process that wakes up and checks various states and takes actions. An example of this is stopping agents that are inactive past the idle limit. In the past this process has been part of the Workspace Manager function, but is now a separate daemon that runs independently. As a result, starting a server brings up this additional server check up daemon (edachkup).

HTTP Listener Daemon

The server has several types of listeners. The HTTP daemon is used for requests using the HTTP protocol. There is one HTTP/TCP Listener process per server.

TCP Listener Daemon

The server has several types of listeners. The TCP daemon is used for requests using the TCP protocol. There is one HTTP/TCP process per server.

FDS Special Services Daemon

The server has several types of special services. The FDS daemon (HLISNK) is used for FOCUS Database requests from agent processes. There is one process per server.

Java Services Daemon

The JSCOM daemon is used for requests using Java. There can be more than one process per server.

Agent Daemons

The actual worker processes for requests are known as agents or tscom3 processes (tscom3 is the actual program name). There is one process per configured agent that is active. Agent processes normally are reused until a recycle point or until they become inactive for a set period. Thus, for a given server, the number of processes associated with agents will change over time.

Gathering Diagnostic Information for Customer Support Services

Diagnostic Service Accelerator (DSA) automates the gathering, packaging, and shipping (using FTP) of trace files and other diagnostic information for Customer Support Services using the edastart -savediag facility. A prompt mode is available for ease-of-use, as well as archiving and FTP posting capabilities. Automatic triggers are available for customer email confirmation. DSA also triggers the logging of the upload to the support case and flags the case as requiring attention, thereby accelerating diagnostics services by automating and closing the communication gap between customer and customer services.

In one step a customer can run a repro, then describe, archive (for example, as a ZIP file), and ship (using FTP) all the required information to Customer Support Services.

Automatic triggers are available for customer email confirmation. DSA also triggers the logging of the upload to the support case and flags the case as requiring attention, thereby accelerating diagnostics services by automating and closing the communication gap between customer and customer services.

Note the following platform-specific information:

- ❑ **On Windows**, users must have a command line version of ZIP installed and on the system path for the archive and send feature to work properly.
- ❑ **All other platforms** use system supplied tools such as tar and pax, and require no additional steps to use the feature.
- ❑ **On IBM i (formerly known as i5/OS) and z/OS**, batch/command files—CMD/CL, JCL, and DCL files, respectively—are available for users who prefer to use this method to control tracing and save diagnostics.

On Windows, the Start menu has equivalent options under the Diagnostics folder.

The files and their equivalent edastart commands are:

File	Command
ITRCON	edastart -tracoon
ITRCOFF	edastart -tracoon
ISAVEDIA	edastart -savediag

The physical execution of these batch files varies by platform. For information about the location and use of supplied batch/command files, see the chapter for your platform in the *TIBCO WebFOCUS® Reporting Server Installation* manual.

Note:

- ❑ The diagnostic process suggests a local location in which to save your files. However, you can supply a custom location if you prefer, as long as you *avoid* reserved system locations including the following: EDHOME; APPROOT; EDAPRFU; SCAROOT; EDACONF bin; EDACONF etc; EDACONF edatemp; EDACONF dfm_dir; EDACONF img_dir; EDACONF catalog; EDACONF web; EDACONF user; any location that is a direct parent in the paths listed above; an ID's HOME; the root of a device (that is, \ or /); c:\windows; c:\Program Files; c:\Documents and Settings; /usr and /usr/bin.
- ❑ In the case of PDS deployment on z/OS, tracing and the edastart -savediag function *must* be done using the ISAVEDIA JCL method described in the z/OS PDS section of the *TIBCO WebFOCUS® Reporting Server Installation* manual (where the diagnostic files go directly to the JES log output of the server) rather than the method described in the remainder of this section.

Procedure: How to Create a Savediag From the Reporting Server browser interface

To run a Savediag, you must be a Server Admin or have administrator privileges.

1. Go to the *Workspace* page. Click *Diagnostics*, then click *Savediag - Report a Bug*.

The Savediag page opens, as shown in the following image.

Savediag

Important:
If your SAVEDIAG includes sensitive content (files or data) do not use options below that perform a post. Instead, send SAVEDIAG to Information Builders using the secure data upload facility at <https://techsupport.informationbuilders.com>.
Contact Customer Support for more information.

Diagnostics Options: Save Diagnostics (savediag) and Post

Copy files from catalog to savediag: Yes No

Note: catalog directory - C:\ibi\srvt77\wfs\catalog

Savediag Directory: C:\ibi\srvt77\wfs\sd032918_144416
Sample: C:\ibi\srvt77\wfs\sd032918_*

Note: You can only change the part of the directory name after the underscore

Comments: [Text Area]

List of Applications: baseapp
Sample: baseapp sql ibisamp app1

Case Number: [Input Field]

Your MedProj ID: [Input Field]

Post ID: [Input Field]

Post Password: [Input Field]

Submit

2. Select from the *Diagnostic Options* drop-down menu.

The options are:

- Save *Diagnostic Options* (savediag) and Post. This is the default.
- Save *Diagnostic Options* (savediag).
- Post existing *Diagnostic Options*.
- Delete existing *Diagnostic Options*.

The available fields will vary, depending on the Diagnostic Option selected.

3. Choose whether to save files from the catalog by selecting the corresponding radio button.
4. For Save options, you can change the part of the *Savediag Directory* field that follows the underscore, which is the time.
5. For Post options, you can enter information in the *Comments* field.
6. For Save options, enter applications in the *List of Applications* field.
7. For Post options, enter the case number you received from Customer Support Services in the *Case Number* field.
8. For Post options, enter the information you received from Customer Support Services in the *Post ID* and *Post Password* fields. This information is required to access the Customer Support Services FTP site.
9. Click *Submit*.

If you selected a Post option, your information is posted on the Customer Support Services FTP site. If not, the information is saved to the *Savediag Directory* folder.

Note: If you choose the *Delete existing Diagnostics* option, select a folder from the *Delete Savediag Directory* drop-down menu and click *Submit*.

Procedure: How to Report TIBCO WebFOCUS Reporting Server Related Problems From the Command Line

In order to provide Customer Support Services with a proper reproduction of your problem, follow these steps:

Note: The use of *edastart* in this procedure assumes that the directory for the *edastart* script has been added to the system search path for your platform (or that another method for locating and using *edastart* has been used). If this is not the case, you must preface *edastart* with its path when you enter it here.

1. Turn traces on with the command *edastart -traceon*. Initially starting the server with traces is strongly recommended.
2. Reproduce the problem while traces are on.
3. If the application abends, write the core dump information to the *edaprint.log* file using one of the following techniques:
 - To create the log from the Reporting Server browser interface menu bar, select *Workspace, Diagnostics, Server Log*, then click *WRITE IBISNAP Information To Current Log*.
 - To create the log manually, issue the *edastart -ibisnap* command.

4. Use the `edastart -savediag` command to gather and send files to the support FTP site as an archive (using the archive method required on your platform). The `edastart -savediag` command includes options for saving and FTPing, saving locally, and FTPing a prior local save. The system prompts as follows:

```
Select an option (type Q to quit for most prompts):
  1. Save diagnostic to a directory and FTP post
  2. Save diagnostic to a directory
  3. FTP post prior diagnostic from a directory

Enter a selection (Default=1) : 1
```

For related information, see [Files Gathered by the edastart -savediag Command](#) on page 799.

Note: Typically, users save and post in one operation (the default), however, you may wish to consider the following reasons for saving and posting as separate steps:

- If you have multiple reproduction variations, you may wish to confirm them all first, then ship them to Customer Support Services.
- If you wish to reproduce multiple separate problems, the two-part option enables you to save each one separately, then ship.
- If the server cannot be shut down (and `edastart -cleardir` cannot be used), the two-part option may enable you to delete unneeded items before shipping.
- For security/privacy reasons, items need to be reviewed and possibly cleansed before shipping.

Procedure: How to Report edastart -t, -x, or -f Problems Outside the TIBCO WebFOCUS Reporting Server Environment

If a problem occurs on the Server environment, attempting to reproduce it in a free-standing environment may facilitate troubleshooting. If you wish to try this strategy, follow these steps.

1. Create a reproduction application directory (preferably called `repro` or `repro_case#`).
2. Place the required files (`*.fex`, `*.mas`, and so on) that are not already in `APP PATH` in the reproduction application directory.
3. Physically switch to the reproduction application directory so that traces and other files go to this known location.

4. Run the desired request with the command:

```
edastart -traceon {-t|-x|-f}
```

where:

```
-t, -x, -f
```

Are three separate trace modes.

5. If the server edatemp directory is not needed to reproduce the problem, shut the server down and issue the edastart -cleardir command so that edatemp is not captured during -savediag processing.

Note: If the server must be kept running 24x7, skip this shutdown step and proceed to step 6.

6. Issue the edastart -savediag command, specifying the repro application directory (and any other needed directories) in order to gather and send files to the support FTP site as an archive (using the archive method required on your platform). The edastart -savediag command includes options for saving locally, saving and FTPing, and FTPing a prior local save. Choose:

1. Save diagnostic in a new directory and post
2. Save diagnostic in a new directory
3. Post diagnostic from an existing directory

```
Enter a selection (Default=1) : 2
```

For related information, see [Files Gathered by the edastart -savediag Command](#) on page 799.

7. Restart the server to resume normal operation.

In the case of an abend (core dump), Customer Support may ask that debuggables be installed (as described in the *TIBCO WebFOCUS® Reporting Server Installation* manual) and that the repro/savediag steps be repeated with IBI_DBGLIB=ON to gather more detailed information. Do not install debuggables or perform these steps unless requested by Customer Support.

Reference: **Files Gathered by the edastart -savediag Command**

- Server log (\$EDACONF/edaprint.log)
- Traces and other contents from the \$EDACONF/edatemp/... directory, including the following:

File	Description
edachkup.trc	Daemon check up facility trace

File	Description
edapdfm.trc	Deferred Listener trace
edaplog.trc	EDAPRINT log daemon trace
fds.trc	FDS trace
jscom*.trc	Java traces
edapth.trc	Main workspace manager trace
edaphtp.trc	HTTP communication traces (listener)
ht000nnn.trc	HTTP communication traces (threads)
edaptcp.trc	TCP/IP communication traces (listener)
ip000nnn.trc	TCP/IP communication traces (threads)
ts000nnn.trc	Agent traces

- Server configuration (\$EDACONF/bin/edaserve.cfg).
- Profiles from \$EDACONF/etc/edasprof.prf and ibi/profiles/*.
- Listener traces (.ht and .hs).
- Deferred execution (.aq and .df).
- Scheduler (.cm and .sh).
- Java Services (.js).
- Complete content of application directories, as specified during -savediag processing.

Example: Sample edastart -savediag Processing With Variations

The following sample includes annotations that indicate what is happening as the process runs.

```
$ bin/edastart -savediag
```

A log of environment information is displayed.

```

.
.
.
Select an option (type Q to quit for most prompts):
  1. Save diagnostic to a directory and FTP post
  2. Save diagnostic to a directory
  3. FTP post prior diagnostic from a directory

Enter a selection (Default=1) : 1

Enter directory path for diagnostic information

(Default=/ul/prod/iadmin/ibi/srv71/ffs/sd122905_110820)

Please supply location or <Enter>:
/ul/prod/iadmin/ibi/srv71/ffs/myrepro

Enter space separated non mapped application names to save in
the diagnostic or ALL for all non mapped application names
Default= baseapp): repro

```

```
12/29/2005 11:08:49
```

```
12/29/2005 11:08:49 Creating savedir =
/ul/prod/iadmin/ibi/srv71/ffs/myrepro
```

- Repro is the application directory for the saved application. myrepro is the directory in which the information generated by -savediag will be saved.
- A log of file copies is displayed. Note that some copy errors are normal since the process is trying to catch files (like those used for NLS set up), which users may or may not have.

```
12/29/2005 11:08:50 Diagnostic Information has been saved in
/ul/prod/iadmin/ibi/srv71/ffs/myrepro
```

- If you selected Option 2, Save diagnostic to a directory, - savediag processing stops here.
- If you selected Option 3, FTP post prior diagnostic from a directory, -savediag processing provides the next prompt.

```
Enter existing directory path with diagnostic information
to post (Default=/ul/prod/iadmin/ibi/srv71/ffs/sd122905_111616)
Please supply location or <Enter>: /ul/prod/iadmin/ibi/srv71/ffs/myrepro
```

It then continues as Option 1 would have.

Please enter Hottrack Case number to post (8 digits or more): 99999999

Please enter user id on cssftp.ibi.com (Default=incmgr):

Please enter comments for the Hottrack Case 99999999.

Enter as many lines as needed, blank line indicates end of comments.

My reproduction for error

12/29/2005 11:09:37 Creating myrepro.tar to post savediag information

Please supply confirmation to proceed with posting (Y/N/Q Default=Y): Y

Note:

- An FTP log of posting activities is displayed.
- As a final step, an HTTP message is sent silently to Customer Support Services. This triggers a confirmation email to the customer and alerts Customer Support Services systems that a –savediag has arrived.

Starting and Stopping the TIBCO WebFOCUS Reporting Server

A server administrator generally handles installation, initial configuration, and on-going management of a server (start, stop, and other maintenance tasks), while developers create applications.

This appendix outlines these basic tasks for a server administrator. If your role is strictly as an end-user or developer, this portion of the documentation is likely irrelevant, unless there is a problem with the server.

In this appendix:

- [Using TIBCO WebFOCUS Reporting Server Management Command Options](#)
- [TIBCO WebFOCUS Reporting Server Environment Variables](#)

Using TIBCO WebFOCUS Reporting Server Management Command Options

The start, stop, and management of a server depends on the platform, but generally a script known as *edastart* is used with parameters to perform required server management tasks (although Windows also has start menu icons, z/OS has JCL procedures that front-end a core JCL, and IBM i has CMD/CL scripts). These control scripts (and applicable parameters) can be incorporated by customer sites to be executed by other means such as another script or a scheduler.

Once the server running, it is often referred to as the workspace, as the server is a collection of processes that coordinate the tasks of managing inbound requests, assigning resources, doing work, and logging events. In this content, these scripts are used without any path qualification (or high-level qualifier on z/OS) that can vary by site. On platforms that support PATH, placing the EDACONF bin directory on the PATH eliminates the need to use a path qualified script name.

The following table outlines specific script locations and names for the supported platforms:

Platform	Location	Name	Notes
Linux, AIX, Solaris	EDACONF bin	edastart	Unix shell script that uses switches to select start, stop, and other functions.

Platform	Location	Name	Notes
Windows	EDACONF bin	edastart.bat	<p>Batch script that uses switches to select start, stop, and other functions. Windows also has:</p> <ul style="list-style-type: none"> <input type="checkbox"/> A Windows start menu group with icons for starting and stopping the server (which are more commonly used). <p>This Start Menu group also has a <i>Diagnostic Functions</i> folder containing an icon for opening a command window in the edastart location, for executing edastart command line operations.</p> <ul style="list-style-type: none"> <input type="checkbox"/> A Windows Service entry with start-up type set to <i>manual</i>. The actual service name matches the workspace name assigned during installation, and the administrator can use the Services application to find the service and configure it, or the NET command to start and stop the service.
IBM i	EDACONF bin	edastart ISTART ISTOP	<p>IFS (QSH) Shell script that uses switches to select start, stop, and other functions (much like any other UNIX environment), but also a matching edastart.sh.</p> <p>IBM i also has CL/CMD native command line source scripts in EDACONF bin to match major tasks (for example, ISTART and ISTOP) as well as compiled versions of these scripts in the installation default library of the ID.</p> <p>These CL and CMD scripts should be rebuilt with unique names if multiple servers are to be run from the same library.</p> <p>If moved to alternate libraries, these CL and CMD scripts should also be rebuilt with the new library names.</p>

Platform	Location	Name	Notes
IBM z/OS HFS Deployment (Uses individual JCL members in an MVS data set for tasks)	<i>hlq.product.DATA</i>	ISTART ISTOP IWAYS IWAYP	<p>Submit this JCL member to start the server as a batch job.</p> <p>Submit this JCL member to stop a server started as a batch job.</p> <p>Submit this JCL member to start the server as a started task*.</p> <p>Submit this JCL member to stop a server started as a started task*.</p> <p>These JCL members front-end the IRUNJCL member for actual execution. In effect, the IRUNJCL is the z/OS equivalent of the edastart script.</p> <p>* Started task JCL must be copied to SYS1.PROCLIB or a JES2 PROCLIB for actual use.</p> <p>To stop the server on the Reporting Server browser interface Workspace page, click <i>Server Actions</i>, then <i>Stop</i>, or in the MVS Console or SDSF Console, use:</p> <p><i>F jobname, -stop</i></p> <p>You can also do an SDSF cancel of the running job</p> <p>For complete information about server installation, configuration, and integration with security packages, see the <i>TIBCO WebFOCUS® Reporting Server Installation</i> manual.</p>

Platform	Location	Name	Notes
IBM z/OS PDS Deployment (Uses individual JCL members in an MVS data set for tasks)	<i>hlq.product.DATA</i>	ISTART IWAYS	<p>Submit this JCL member to start the server as a batch job.</p> <p>Submit this JCL member to start the server as a started task*.</p> <p>These JCL members front-end the IRUNJCL member for actual execution. In effect, the IRUNJCL is the z/OS equivalent of the edastart script.</p> <p>* Started task JCL must be copied to SYS1.PROCLIB or a JES2 PROCLIB for actual use.</p> <p>You can also stop the server on the Reporting Server browser interface Workspace page by clicking <i>Server Actions</i>, then <i>Stop</i>, or in the MVS Console or SDSF Console, using:</p> <p><code>F jobname, -stop</code></p> <p>You can also do an SDSF cancel of the running job</p> <p>For complete information about server installation, configuration, and integration with security packages, see the <i>TIBCO WebFOCUS® Reporting Server Installation</i> manual.</p>

Reference: Options for Starting and Stopping the TIBCO WebFOCUS Reporting Server

On Windows, there is a Windows start menu icon group for starting and stopping the server. Additionally, there is a Reporting Server browser interface icon to access it without having to type the URL into the browser. There are other icons, and folders with icons, that allow direct access to editing configuration files and other tasks, such as turning on traces or opening a command line session, and to use the edastart command options described in this section.

The following table lists the edastart options for starting and stopping the server.

Command and Option Flag	Purpose	Notes
<p><code>edastart</code></p> <p>(no option flag)</p>	<p>Starts the server and also displays the server log (edaprint) to the screen as an interactive log on non-IBM z/OS platforms.</p> <p>Windows has a <i>Start Security ON</i> icon in the Windows start menu icon group for this task. Additionally, the server may be started by starting the associated Windows Service, but logging is strictly to the server log (edaprint).</p> <p>IBM z/OS (HFS or PDS Deployment) has an equivalent <i>hlq.product.DATA(ISTART)</i> JCL member for this task, but the log (edaprint) becomes one of the SDSF logs, and operator commands can be used to interact with the job.</p>	<p>Interactive logging can be interrupted using the Ctrl+C character combination, except on IBM i (which uses SYSREQ and 2) and z/OS (which does not support interactive logging). Most other edastart options can be passed (for example, status), which displays output, and then logging resumes.</p> <p>Interrupting interactive logging is typically used to turn tracing on or off (-traceon/-traceoff), but most options are valid.</p> <p>To exit interactive logging, interrupt the interactive log and use the -quit option. The interactive log may be resumed, with edastart -console.</p> <p>When you issue an option in interactive logging mode, the leading dash (-) may be left off the option.</p>
<p><code>edastart -start</code></p>	<p>Starts the server and returns user control to the command line immediately. An initial start message will display, but all other logging goes to the EDACONF edaprint.log file.</p> <p>IBM z/OS (HFS or PDS Deployment) has an equivalent <i>hlq.product.DATA(ISTART)</i> JCL member for this task.</p>	<p>Windows has no equivalent start menu icon. The server may be started by starting the associated Windows Service, where logging goes to the server log (edaprint).</p>

Command and Option Flag	Purpose	Notes
<code>edastart -sstart n</code>	Starts the server, but waits <i>n</i> seconds to display additional sub-service start confirmation messages before returning user control to the command line. Control passes back to the command line when all services finish starting or when the timeout is reached.	<p>Warning: Too short a wait period can give a false impression of startup failure, as the fail messages indicate a failure to start within the period. Sites must adjust the wait period to be reasonable for the load of the machine. A period of 15 seconds seems to be a reasonable number for most machines we have encountered.</p> <p>This feature does not apply to IBM z/OS (HFS or PDS Deployment), nor is there an equivalent Windows start menu icon.</p>
<code>edastart -stop</code>	Stops the server. Sends a stop signal to the server, but does not wait for the final stop message before returning control to the command line.	<p>Windows has a Stop icon in the Windows start menu icon group for this task. The server can also be stopped by stopping the associated Windows Service.</p> <p>IBM z/OS (HFS Deployment) has an equivalent <code>hlq.product.DATA(ISTOP) JCL</code> member for this task, or use the following command in the MVS Console or in SDSF.</p> <p><code>F jobname, -stop</code></p> <p>For IBM z/OS PDS Deployment, the equivalent is only to use the MVS Console or SDSF.</p>

Command and Option Flag	Purpose	Notes
<code>edastart -sstop n</code>	Stops the server, but waits n seconds to display the final stop confirmation message before returning user control to the command line.	<p>Warning: Too short a wait period can give a false impression of stop failure, as the fail messages indicate a failure to stop within the period. Sites must adjust the wait period to be reasonable for the load of the machine. A period of 15 seconds seems to be a reasonable number for most machines we have encountered.</p> <p>This feature does not apply to IBM z/OS (HFS or PDS Deployment).</p>

Reference: Other Options

The following table lists commonly used edastart options for performing tasks other than starting or stopping the server.

Command and Options	Purpose	Notes
<code>edastart -?</code> <code>edastart -??</code> <code>edastart -?s</code>	<p>Displays basic workspace control options by area:</p> <ul style="list-style-type: none"> <input type="checkbox"/> edastart -? lists options for Basic Workspace Control. <input type="checkbox"/> edastart -?? lists options for Debugging and Internal Utilities. <input type="checkbox"/> edastart -?s lists options for Support and Administration. 	<p>IBM z/OS (HFS and PDS Deployments) has no equivalent JCL member for this task. Use the following command in the MVS Console or the SDSF Console.</p> <p><code>F jobname, -?</code></p>
<code>edastart -status</code>	Displays whether the server is up or down.	<p>IBM z/OS (HFS and PDS Deployments) has no equivalent JCL member for this task. Use the following command in the MVS Console or the SDSF Console.</p> <p><code>F jobname, -status</code></p>

Command and Options	Purpose	Notes
<code>edastart -quit</code>	When the server is started in interactive log mode and the session is interrupted, the quit option exits the user from interactive logging mode and returns to the operating system command line.	The -quit option exists solely for the purpose of exiting the interactive console display. As an edastart command line option, it is ignored.
<code>edastart -console</code>	Resumes the current interactive log session of a running server, after a quit has been executed.	IBM z/OS (HFS or PDS Deployment) has no equivalent JCL member for this task, use the SDSF Console to show and work with the job log.
<code>edastart -show</code>	Displays workspace status.	<p>IBM z/OS (HFS Deployment) has an equivalent <code>hlq.product.DATA(ISHOW)</code> JCL member for this task.</p> <p>There is no equivalent job for PDS Deployment. However, for both deployment modes, you can use the following command on the MVS Console or the SDSF Console.</p> <p><code>F jobname, -show</code></p>
<code>edastart -showlog</code>	Displays the current server log.	<p>IBM z/OS (HFS and PDS Deployments) has no equivalent JCL member for this task. Use the following command in the MVS Console or the SDSF Console.</p> <p><code>F jobname, -showlog</code></p>
<code>edastart -restart</code>	Restarts the server.	<p>IBM z/OS (HFS and PDS Deployments) has no equivalent JCL member for this task. Use the following command in the MVS Console or the SDSF Console.</p> <p><code>F jobname, -restart</code></p>

Command and Options	Purpose	Notes
<code>edastart -clear</code>	Clears inter-process shared memory controls. When a server is abruptly stopped in an uncontrolled manner, such as kill -9, crash, or power failure, occasionally resources will be locked and prevent server restart. The -clear option will, in most cases, clean these resources, allowing the start to proceed. Otherwise, the processes must be manually killed, followed by issuing <code>edastart -clear</code> .	<p>IBM z/OS HFS Deployment has an equivalent <code>hlq.product.DATA(ICLEAR)</code> JCL member for this task or use the following command in the MVS Console or SDSF Console.</p> <pre>F jobname, -clear</pre> <p>This feature does not apply to IBM z/OS PDS Deployment, as it does not use shared memory.</p>
<code>edastart -cleardir</code>	Cleans up agent run-time directory and files in EDACONF.	<p>IBM z/OS HFS Deployment has an equivalent <code>hlq.product.DATA(ICLRDIR)</code> JCL member for this task.</p> <p>This feature does not apply to IBM z/OS PDS Deployment because it uses a collection of temporary z/OS files for <code>edatemp</code>, which disappear automatically on server shutdown.</p> <p>Warning: Do not use this feature if your EDACONF location contains files that you do not want deleted, such as RM log files.</p>

Command and Options	Purpose	Notes
<pre>edastart -traceon</pre>	<p>Starts server tracing.</p> <p>The -traceon option may be used in conjunction with any server start command. If the server is not already running, the -traceon option will also start the server.</p>	<p>The Reporting Server browser interface User menu also has an Enable/Disable Traces option.</p> <p>Windows has an equivalent <i>Start ... with Traces</i> start menu icon in the Diagnostic Functions folder.</p> <p>IBM z/OS (HFS and PDS Deployment) has an equivalent <code>hlq.product.DATA(ITRCON)</code> JCL member for this task, or use the following command in the MVS Console or SDSF Console.</p> <pre>F jobname, -traceon</pre>
<pre>edastart -traceoff</pre>	<p>Stops server tracing.</p> <p>The -traceoff option may be used in conjunction with any server command.</p>	<p>The Reporting Server browser interface User menu also has an Enable/Disable Traces option.</p> <p>IBM z/OS HFS Deployment has an equivalent <code>hlq.product.DATA(ITRCOFF)</code> JCL member for this task</p> <p>For both IBM z/OS HFS and PDS Deployments, you can use the following command in the MVS Console or SDSF Console.</p> <pre>F jobname, -traceoff</pre>

Command and Options	Purpose	Notes
<code>edastart -savediag</code>	<p>Saves the server state to a directory, compresses it and offers to post it to the software support case management system (an already open case is required for posting).</p> <p>This command is usually used in conjunction with running traces for problem reproduction and for transmitting the reproduction and the traces to software support.</p>	<p>Windows has an equivalent <i>Save Diagnostic</i> start menu icon in the Diagnostic Functions folder.</p> <p>IBM z/OS (HFS and PDS Deployments) has an equivalent <code>hlq.product.DATA(ISAVEDIA) JCL</code> member for this task, or use the following command in the MVS Console or SDSF Console.</p> <p><code>F jobname, -savediag</code></p>
<code>edastart -nohs</code>	Blocks Hyperstage from starting at server startup.	<p>Note: The Hyperstage feature does not apply to all platforms, nor is Hyperstage included in initial installations, by default, as it is deprecated.</p>

TIBCO WebFOCUS Reporting Server Environment Variables

Several environment variables can be used to affect direct server operations. Additionally, certain environment variables must be set in order to configure or use some DBMS adapters.

Depending on your platform (Windows, UNIX, Linux, IBM i, or IBM z/OS), you may want to set these control variables at a machine level, a user ID level, or within a calling script. The server also supports an environment configuration file (`edaenv.cfg`) where `name=value` environment variable pairs (one per line) can be placed and activated into the environment of the process at start up.

You can add `name=value` environment variable pairs to the `edaenv.cfg` file in several ways, but the server must be restarted in order to take advantage of any new settings. Variable values in the environment configuration file must be plain strings (variables within values are not allowed. For example, `X=$HOME/abc` is not allowed). Spaces within a value string is allowed, and enclosing the value in single or double quotation marks is optional.

You can use the following methods to add, change, or delete value pairs in the `edaenv.cfg` file.

- ❑ All platforms can use the server Reporting Server browser interface configuration page to add references (one pair per line). For access to the configuration page, sign into the Reporting Server browser interface and navigate to the Workspace page. Expand *Configuration Files*, expand *Miscellaneous*, right-click *Environment - edaenv.cfg*, and click *Edit* (which will open an editor environment). Add your environment variable name=value pairs, each on a separate line, save the changes, and click *Save and Restart Server*.
- ❑ On Windows, the menu icon group has a *Configure Functions* folder, and within the folder is an *Edit Environment Settings File - EDAENV.CFG* item to manually edit the file (add your environment variable name=value pairs, each on a separate line, then save, and restart the server), or use file explorer to find and edit the `EDACONF bin\edaenv.cfg` file in Notepad or your favorite text editor.
- ❑ On UNIX, Linux, AIX, IBM i, or Windows, edit the `EDACONF bin edaenv.cfg` file (add your environment variable name=value pairs, each on a separate line, then save and restart server).
- ❑ On IBM z/OS edit the `hlq.product.DATA(EDAENV)` file (add your environment variable name=value pairs, each on a separate line, then save, and restart the server).

Note: The server normally does not ship with an `edaenv.cfg` file, so it is normal to get "new file" warnings from the editor you may be using.

The following variables control basic server operation:

EDATEMP

Redefines where the server `EDACONF *.log` files and `edatemp` directory are created.

EDAIPC

Defines an alternate local disk location for server pipe and fifo creation when the `EDACONF` configuration is on an NFS-mounted disk or `EDATEMP` is declared and pointing to an NFS-mounted disk. This only applies to Linux/UNIX systems. For more information on setting `EDAIPC`, see [EDATEMP and NFS-Mounted Disks](#) on page 833.

EDAPROF

Redefines where the server profile, `edasprof.prf`, is located. The default location is `EDACONF etc`. The value is the `edasprof.prf` location (and is *not* an alternate path and file name).

JAVA_HOME

Is the location of a Java JRE installation. This variable is used to find the Java JVM DLL for the Java Listener. On Windows, the registry is checked for an installed copy and, if found, it is used, but setting the variable allows even the registry to be overridden.

JDK_HOME

Is the location of a Java JDK installation. This variable is used to find the Java JVM DLL for the Java Listener. On Windows, the registry is checked for an installed copy and, if found, it is used, but setting the variable allows even the registry to be overridden. If both JAVA_HOME and JDK_HOME are set, the JDK_HOME value is used.

_JAVA_OPTIONS

Enables you to set additional Java JVM start-up parameters for the Java Listener.

Use of _JAVA_OPTIONS is not explicitly recommended because it triggers JVM warning messages, such as *Picked up _JAVA_OPTIONS: ...*, which in turn cause adapter bulk load sessions to think a real error has occurred (which causes a rollback request to be issued, so the data is never loaded).

If adapters using Java-based bulk load facilities are not in use, then use of _JAVA_OPTIONS is not problematic. However, when Java-based bulk load facilities are in use, setting the desired JVM options on the Reporting Server browser interface Workspace *Java Services Properties* page (JVM Settings section) will set values in a way that does not trigger these problematic messages.

EDACS3

Redefines the location and name of the server communications file (odin.cfg).

EDAPRFU

Redefines the location where the server admin.cfg file and user profiles reside.

FOCPRI

Is a legacy method for specifying the OFFLINE printing tool when OFFLINE CLOSE is issued. The default is *lp \ \$1*, and this setting allows you to add a parameter, such as a printer location.

The FOCPRINT feature is considered deprecated, as the FILEDEF command (DYNAM on z/OS PDS Deployment) has, for many years, supported setting OFFLINE PRINTER device specifications and can be used within normal application code. For more information, see [OFFLINE Printing](#) on page 561.

EDAPRIV

Is an EDASHOME/EDACONF overlay, that is, an alternate location containing one or more directories (possibly containing additional subdirectories) that model the directories of EDASHOME and EDACONF.

Test versions of files may be placed in their respective EDAPRIV locations, and the server, upon start-up, will see and use the alternate files. This feature will also act as though EDATEMP has been declared (so it will contain *.log and edatemp files), unless EDATEMP is explicitly also declared.

This is a testing feature and should not be used for production purposes.

EDAPRIV_PORT

Is the server port range offset number.

If the original EDACONF port range of a server starts at 8100, setting EDAPRIV_PORT to an offset of 100 will cause the port range to start at 8200. Therefore, the port numbers will not conflict with another configuration with the same port range (if it is running at the same time).

HOMEAPPS

Overrides the homeapps location in the server configuration.

WFRS_NOHS

Setting a value of TRUE blocks Hyperstage from starting at server startup. The Hyperstage feature does not apply to all platforms, nor is Hyperstage included in initial installations, by default, as it is deprecated.

IBI_LIBPATH

IBI_LIBPATH is a generic variable that notifies the server start-up process that there are additional directories to be added internally to the library path of the server, regardless of platform. Multiple directories are allowed in the string value, but use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

The general purpose is for making third-party DLL driver locations accessible to the server (typically for data adapters).

While some platforms allow other system variables to adjust library path (such as PATH on Windows, LIBPATH on AIX, STEPLIB on z/OS, and LD_LIBRARY_PATH on Linux and Solaris), secured (tscom300 SUID) servers will ignore these methods as an operating system security feature. Therefore, IBI_LIBPATH should be used (even if the server is unsecured, for consistency).

The server also recognizes several new and individual variables as "add to DLL path" variables for declaring product-specific third-party DLL API locations, so the declarations can be done in an organized manner from obviously named variables (see IBI_LOAD_*_FROM variables in this list). While the prior path load methods continue to work, it is preferable to use the new IBI_LOAD_*_FROM variables because of the organization it gives to the environment.

IBI_LOAD_SSL_FROM*

Normally, SSL and CRYPTO DLLs are installed by system administrators as a secondary package in system default locations that would normally be on the standard system library path. It is the responsibility of the system administrator to make sure the DLLs are available, if the server application code makes use of these functions.

Occasionally, there may be reason to use an alternate SSL and CRYPTO version in an alternate location. The IBI_LOAD_SSL_FROM environment variable enables you to set a load directory where the applicable SSL and CRYPTO DLLs exist. They must match the bit size of the server (usually 64bit) and will be added to the library path at server runtime.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

IBI_LOAD_MSSQL_ODBC_DRIVER_FROM*

Is the Microsoft SQL Server and Azure ODBC driver location.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

IBI_LOAD_ORACLE_CLIENT_FROM*

Is the Oracle Client DLL library location.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

IBI_LOAD_REDSHIFT_ODBC_DRIVER_FROM*

Is the Redshift ODBC driver location.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

IBI_LOAD_EXASOL_ODBC_DRIVER_FROM*

Is the EXASolution ODBC driver location.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

IBI_LOAD_DB2_CLI_CLIENT_FROM*

Is the Db2 CLI Client DLL library location.

Normally this variable only needs to point to one directory. However, if multiple directories are needed, use a semi-colon (;) as the separator on Windows and a colon (:) on all other platforms.

* Prior releases recommended the use of the LD_LIBRARY_PATH, LIBPATH, STEPLIB, or IBI_LIBPATH variable to adjust the library path for loading third-party DLLs, depending on whether system security was in use. While these methods continue to work, they are deprecated in favor of using the IBI_LOAD_*_FROM variables.

Unicode Support

Unicode is a universal character encoding standard that assigns a code to every character and symbol in every language in the world. Since no other encoding standard supports all languages, Unicode is the only encoding standard that ensures that you can retrieve or combine data using any combination of languages.

In this appendix:

- [Unicode and the TIBCO WebFOCUS Reporting Server](#)
 - [Accessing Unicode Data](#)
 - [Selecting, Reformatting, and Manipulating Characters](#)
 - [Sort Order Under Unicode](#)
 - [Added Unicode Support for Master Files, Data Files, and Application Directory Names](#)
 - [Unicode PDF Output](#)
-

Unicode and the TIBCO WebFOCUS Reporting Server

The WebFOCUS Reporting Server supports a Unicode Transformation Format (UTF) called UTF-8 in ASCII environments, and UTF-EBCDIC in EBCDIC environments:

- For ASCII**, the UTF-8 encoding standard assigns each character of each language a code that can be from 1-3 bytes long. The codes assigned to European characters are 1 or 2 bytes long, Middle Eastern characters are 2 bytes long, and those assigned to Asian characters are 3 bytes long. This standard is compatible with ASCII format because the first 128 UTF-8 codes have the same 1-byte representation as the corresponding ASCII codes.
- For EBCDIC**, the UTF-EBCDIC encoding standard assigns each character a code that can be from 1-4 bytes long. EBCDIC characters, including C1 control characters, have the same 1-byte representation in UTF-EBCDIC.

In non-Unicode single-byte encoding standards, such as ASCII, each character is assigned a code that is 1 byte long, limiting the number of characters that can be represented by the standard. When using those standards, it became common to equate a character with a byte of storage. If you had a string of 10 characters, the amount of storage needed was 10 bytes, and many character manipulation routines expected character string lengths to be specified as a number of bytes.

With Unicode encoding, bytes and characters are no longer equated: characters are represented internally by a varying number of bytes, depending on the character. If you configure the server for Unicode, you define the length of strings and alphanumeric fields in terms of characters, not bytes. This simplifies specifying string and field lengths. Each character is represented internally by up to 3 bytes, and the server automatically adjusts for the actual storage length. In reports, each character displays in a report column using one space, regardless of how many bytes it takes up in memory. This character-based processing mode employed for Unicode environments is called *character semantics*. The non-Unicode mode is called *byte semantics*.

Procedures that had been developed using byte semantics will continue to work when deployed in a Unicode environment, without adjustment, in most cases.

To compress trailing blanks and display columns as the width of the largest actual data value, issue the SET SQUEEZE=ON command. For more information on the SQUEEZE parameter, see the chapter on Formatting Report Data in the *Creating Reports With TIBCO WebFOCUS® Language* manual.

The main benefit of the new system is the ability to have multiple languages (both European and Asian) in the following WebFOCUS and Dialogue Manager objects:

- Titles, descriptions, and names in synonyms.
- Headings and prompts in procedures.
- Data for all supported adapters (for example, SAP BW, SAP R/3-ECC, Oracle, Db2, Sybase ASE, Sybase IQ, Teradata, MySQL, Web Services, Fixed files). For more information, see [Accessing Unicode Data](#) on page 821.

Procedure: How to Configure the Server for Unicode

To configure the server for UTF-8 or UTF-EBCDIC character encoding:

1. Go to the *Workspace* page.
2. Click *Settings*, point to *LOCALE*, and click *Configuration Wizard*.

The NLS Configuration Wizard pane opens.

3. Select *65001 - Unicode (UTF-8)* or *65002 - Unicode (UTF-EBCDIC)* from the CODE_PAGE drop-down list.

The server will be configured for Unicode once you save this configuration.

Accessing Unicode Data

To find out which adapters support Unicode:

1. Go to the *Get Data* page on the Reporting Server browser interface.
2. In Simple Mode, click *Information and Diagnostics* or, in Advanced Mode, click *Options*, then click *Complete Adapter List*.

This opens a list of all adapters. There is a Unicode column. The value Opt in this column indicates that the adapter supports Unicode.

Relational adapters in a Unicode environment assume that the DBMS returns character data to the server already converted to Unicode. The relational adapters convert data to the correct DBMS API when writing to a relational data source (for example, Oracle to UTF-8, Microsoft SQL Server to UTF-16, and Db2 on MVS to UTF-EBCDIC).

XML-based adapters (the Adapter for XML, and the Adapter for XBRL) obtain the code page from the XML declaration of the processed XML document. For more information, see <http://www.w3.org/TR/REC-xml#sec-prolog-dtd>.

The Adapter for Web Services generates SOAP requests using the UTF-8 code page.

Reference: Unicode Considerations for Oracle

The adapter supports Unicode data in Oracle release 10g or higher databases that have been configured with the NLS_CHARACTERSET parameter set to UTF8. You must set the NLS_LANG environment variable in the edastart file, in a separate shell file, in a database profile, or in a user profile.

Set NLS_LANG using the following syntax

```
NLS_LANG = language_territory.characterset
```

where:

language

Is the selected language.

territory

Is the name of the country associated with the selected language.

characteraset

Is the value of the NLS_CHARACTERSET variable that is set in the Oracle database. For Unicode, this is always UTF8.

For example, for American English UTF-8, you would use the following setting:

```
NLS_LANG=American_America.UTF8
```

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: Unicode Considerations for Db2

The server supports Db2 databases, version 8 and higher. To prepare the Db2 environment for Unicode on:

- ❑ **Windows**, the database must have been created with the option CODESET UTF-8, and you must add the following variable to the environment using Windows or in the edastart file:

```
DB2CODEPAGE=1208
```

- ❑ **UNIX**, the database must have been created with the option CODESET UTF-8, and you must set the LANG and NLS_LANG environment variables in the edastart file or in a separate shell file.

For example, for American English, you would export the following variables:

```
export LANG=EN_US.UTF-8
export NLS_LANG=American_America.UTF8
```

- ❑ **z/OS**, the database must have been created with the CCSID UNICODE option, and you must ensure that the DSNAINI environment variable points to a configuration file containing the following specification:

```
CURRENTAPPENSCH=UNICODE
```

The adapter supports Unicode only with the CLI interface.

In a Unicode environment, the Adapter for Db2 requires a BIND command for PREPARE/EXECUTE logic using parameter markers.

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

- ❑ **IBM i**, Unicode is only supported in Db2 CLI mode. Db2 allows column-by-column specification of CCSID information at CREATE TABLE time and the columns may or may not be explicitly Unicode 1208 and 1200, however, all CCSIDs are transcoded to the server as CSID 1208 UTF8. Thus, any existing table may be used in a Unicode configuration regardless of its underlying CCSID specifications.

Reference: **Unicode Considerations for Sybase ASE**

The adapter supports Unicode data in Sybase ASE version 15.0 and higher databases that have been created with the CHARACTER SET option set to UTF-8. You must set the LANG and NLS_LANG environment variables in the edastart file or in a separate shell file before starting the server.

For example, for American English, you would export the following variables:

```
export LANG=EN_US.UTF-8
export NLS_LANG=American_America.UTF8
```

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: **Unicode Considerations for Sybase IQ**

Beginning with Sybase IQ version 12.7, the adapter supports Unicode data in Sybase IQ databases that have been created with the UTF-8 character set. You must set the LANG and NLS_LANG environment variables in the edastart file or in a separate shell file before starting the server.

For example, for American English, you would export the following variables:

```
export LANG=EN_US.UTF-8
export NLS_LANG=American_America.UTF8
```

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: **Unicode Considerations for Microsoft SQL Server**

The adapter, using the OLE DB interface, supports Unicode data stored in NCHAR and NVARCHAR fields (where N stands for national). N columns can support data of any language or combination of languages.

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: Unicode Considerations for Teradata (CLI)

The Adapter for Teradata (CLI) supports Unicode UTF-8 format if:

- The Teradata CLI client components are part of release TTU8.0 or higher.
- The Teradata database is release V2R6.0 or higher and appropriate language support was enabled during the sysinit process.

Contact your database administrator (DBA) to determine whether international language support has been enabled in your Teradata system and/or consult the Teradata documentation for details about International Character Set support.

Note that, at the present time, when Unicode is enabled the length of a Teradata Column Name and/or TITLE cannot exceed 21 characters (bytes).

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: Unicode Considerations for MySQL

The Adapter for MySQL is implemented using JDBC. This implementation supports Unicode data stored in character fields with CHARACTER SET set to UTF-8.

You must set the LANG environment variable in the edastart file or in a separate shell file before you start the server. For example, for American English you would export the following variable:

```
export LANG=EN_US.UTF-8
```

For information about data type support, see [Relational Adapter Data Type Support for Unicode](#) on page 825.

Reference: Relational Adapter Data Type Support for Unicode

In Unicode databases the information in CHAR(*n*) columns is stored in a UTF-8 encoding scheme. Most RDBMS Unicode columns of CHAR type specify length in bytes, not characters. The B-modifier in the Actual format denotes that a character column with a fixed byte length might contain a varying number of UTF-8 characters. This is reflected in the AnV Usage format.

DBMS	Column Type	Usage	Actual*
Db2	CHAR(<i>n</i>)	AnV	AnB
	GRAPHIC(<i>n</i>)	An	An
	VARCHAR(<i>n</i>)	AnV	AnVB
	VARGRAPHIC(<i>n</i>)	AnV	AnV
Microsoft SQL Server	CHAR(<i>n</i>) single byte code page	An	An
	CHAR(<i>n</i>) double byte code page	AnV	AnV
	NCHAR(<i>n</i>)	An	An
	VARCHAR(<i>n</i>)	AnV	AnV
	NVARCHAR(<i>n</i>)	AnV	AnV
MySQL	CHAR(<i>n</i>)	An	An
	VARCHAR (<i>n</i>)	AnV	AnV
Oracle	CHAR(<i>n</i> CHAR)	An	An
	CHAR(<i>n</i> BYTE)	AnV	AnB
	NCHAR(<i>n</i>)	An	An
	VARCHAR(<i>n</i> CHAR)	AnV	AnV
	VARCHAR(<i>n</i> BYTE)	AnV	AnVB
	NVARCHAR(<i>n</i>)	AnV	AnV

DBMS	Column Type	Usage	Actual*
Sybase ASE	CHAR(<i>n</i>)	<i>An</i>	<i>AnB</i>
	UNICHAR(<i>n</i>)	<i>An</i>	<i>An</i>
	VARCHAR(<i>n</i>)	<i>AnV</i>	<i>AnVB</i>
	UNIVARCHAR(<i>n</i>)	<i>AnV</i>	<i>AnV</i>
Sybase IQ **	CHAR(<i>n</i>)	<i>An</i>	<i>AnB</i>
	VARCHAR(<i>n</i>)	<i>AnV</i>	<i>AnVB</i>
Teradata	CHAR(<i>n</i>)	<i>An</i>	<i>An</i>
	VARCHAR (<i>n</i>)	<i>AnV</i>	<i>AnV</i>

* Note that on EBCDIC platform(s) the ACTUAL size for a B-suffixed format is increased 1.5 times to accommodate the expansion when converting from UTF-8 to UTF-EBCDIC. For example, on MVS the synonym created for a Db2 CHAR(10) column contains the following: USAGE=A10, ACTUAL=A15B.

****Note the following limitation for Sybase IQ:** You cannot use the HOLD FORMAT SYBASE command with the Unicode implementation of Sybase IQ since this command depends on the availability of UNICHAR and UNIVARCHAR data types, which are not supported by Sybase IQ.

Reference: Unicode Considerations for SAP BW and SAP R/3-ECC

SAP uses UTF-16 encoding in its Unicode system. The server uses UTF-8 and handles all conversions between the two encoding schemes. The server may not need to be configured for Unicode when accessing the SAP Unicode system.

NLS settings for the server must be configured in such a way that the server code page can handle the list of chosen languages. For example, ISO 8859-1 can accommodate most Western European languages. The 8859 family can handle character specifics with the lower set almost being mapped to US ASCII. Therefore, with 8859-1 one could request English, German, French, and Spanish. When a character set requires a code page that takes more than one byte per character (for example, many Asian languages), the only choice for the server is 65001 (UTF-8).

The adapters provide access to Unicode SAP BW and SAP ECC systems, respectively. This extends support of data and metadata in multiple languages to the server, consistent with support by the SAP server. A synonym can be created using one or more languages. Those languages will be used to create titles and descriptions.

- ❑ For SAP BW, the userid and password in the `sapserv.cfg` file must be able to connect to SAP BW using the enumeration of desired languages.
- ❑ For SAP R/3-ECC, the server logon language is used to retrieve all languages.

Reference: **Unicode Considerations for Fixed-Format Sequential Files**

When retrieving a fixed-format sequential file, the server attempts to determine the code page the file was meant to be retrieved with by checking the Master File CODEPAGE attribute. If the Master File does not contain the CODEPAGE attribute, the server code page is used to read the file.

If you use the Data Migrator desktop interface to generate a data flow that creates a fixed-format sequential file, you can specify a code page in the Data Migrator desktop interface. Data Migrator will then create the fixed-format file in a way that can be read by the server when that server has been configured for the specified code page.

In a Unicode configuration, HOLD files in BINARY and ALPHA formats are created using UTF-8 conversion, which assigns each character three bytes of storage in ASCII environments or four bytes in EBCDIC environments. Fields defined in the Master File using the data type A in both the USAGE and ACTUAL attributes are described in terms of characters. Fields defined using any other combination of USAGE and ACTUAL attribute values are described in terms of bytes.

To force a field in a fixed-format sequential file to be described in terms of bytes, add B to the end of the ACTUAL attribute. For example, to specify that a field is stored in 10 bytes, you would specify:

```
ACTUAL=A10B
```

The adapter will then read the specified number of bytes from the record and convert their contents to the number of characters specified by the file code page.

Regardless of how much storage a character occupies, it occupies only one space on a report, as always.

Selecting, Reformatting, and Manipulating Characters

In character semantics mode, selection tests against a mask are automatically adjusted to work with characters rather than bytes. Formats assigned by reformatting a field in a request or by defining a temporary field are interpreted in terms of characters. Character functions interpret all lengths in terms of characters.

Example: Defining a Virtual Field

Consider the following DEFINE in the Master File for the EMPLOYEE data source:

```
DEFINE FIRST_ABBREV/A5 WITH FIRST_NAME = EDIT(FIRST_NAME, '99999$$$$$');$
```

In character semantics mode, format A5 is interpreted as five characters (up to 15 bytes on ASCII platforms, up to 20 bytes on EBCDIC platforms), and the comparison is performed based on this number of bytes. In byte semantics mode, format A5 is interpreted as five bytes, and the comparison is performed based on five bytes. In either case, the correct characters are compared and extracted.

Example: Reformatting a Field

Consider the following PRINT command:

```
PRINT FIELD1/A10
```

In character semantics mode, format A10 is interpreted as 10 characters (up to 30 bytes), meaning that up to 30 bytes must be retrieved when this field is referenced. In byte semantics mode, format A10 means that 10 bytes will be retrieved. In either case, the field displays as 10 characters that take up 10 spaces on the report output.

Reference: Character Functions That Support Character Semantics

In character semantics mode, all character manipulation functions interpret lengths in terms of characters. The following functions operate on alphanumeric strings in character semantics mode when Unicode is configured:

String manipulation and extraction functions.

GETTOK, OVRLAY, PARAG, REVERSE, SQUEEZ, STRIP, SUBSTR, SUBSTV, TRIM, TRIMV

Justification functions.

CTRFLD, LJUST, RJUST

Length and position functions.

ARGLEN, LENV, POSIT, POSITV

Format conversion functions.

EDIT

Decoding, comparison, and editing functions.

CHKFMT, EDIT, DECODE, SOUNDEX

String replacement functions.

CTRAN, HEXBYT, BYTVAL (see notes below), STRREP

Case translation functions.

LCWORD, LOCASE, LOCASV, UPCASE, UPCASV

Note: The HEXBYT, BYTVAL, and CTRAN functions have been extended to handle multibyte characters in Unicode configurations. These functions use or produce numeric values to represent characters. In Unicode configurations, they use or produce values in the range:

- 0 to 255 for 1-byte characters
- 256 to 65535 for 2-byte characters
- 65536 to 16777215 for 3-byte characters
- 16777216 to 4294967295 for 4-byte characters (primarily for EBCDIC)

To find the numeric value corresponding to a given character, find its hexadecimal code and convert to decimal with a hex calculator such as the Windows XP Calculator program. (Make sure to use the UTF-8 or UTF-EBCDIC code, not the Unicode code point, which would be the UTF-16 value.)

For example, assume you would like to create a variable of format A1 containing the euro sign. The euro sign in UTF-8 is, in hex, E282AC. Converting this to decimal gives 14849492. Thus, the proper DEFINE or COMPUTE would be:

```
EUROSIGN/A1 = HEXBYT(14849492, 'A1');
```

If you are creating a FOCEXEC with a UTF-8 compliant editor, you can also get the value of the euro sign in this way:

```
EUROVAL/I8 = BYTVAL('€', 'I8');
```

The CTRAN function replaces all occurrences of a character in a string with another character, given the decimal values that represent the hexadecimal codes for the two characters. Traditionally, this technique was used to replace characters that were difficult to input directly. Decimal values of characters can be complicated to determine. Therefore, if you want to replace characters or character strings that you can input directly using a UTF-compliant text editor, use the STRREP string replacement function.

The following translates all of the euro signs in a 40-character UTF-8 field to pound sterling signs (£ = C2A3 or 49827):

```
NEWFLD/A40 = CTRAN(40, OLDFLD, EUROVAL, 49827, 'A40');
```

Sort Order Under Unicode

Sort order is based on the binary values assigned to the characters. When the server is configured for Unicode, the sort order is based on the Unicode encoding standard. If ascending values of the codes correspond to the alphabetical order of the letters in the language being used, a report can be sorted in alphabetical order. This is entirely dependent on the encoding standard and its mapping of codes to letters. In many, but not all cases, the encoding standard assigns codes in the alphabetical order of each language.

For example, Ukrainian added a new letter (Cyrillic capital letter ghe with upturn) to its alphabet after the UTF-8 coding specification had already been set. This letter was not assigned a code that sorts it alphabetically, either in Unicode or in code page 1251 (used for Ukrainian). It sorts differently and incorrectly using either encoding scheme.

- ❑ With code page 1251, this letter sorts as the first letter on the report output.
- ❑ With UTF-8, this letter sorts as the last letter on the report output.

To determine whether a language sorts alphabetically, you can examine the hexadecimal codes assigned to its letters on the code page you are using and check whether ascending hexadecimal codes match the alphabetical order.

Added Unicode Support for Master Files, Data Files, and Application Directory Names

Support for UTF-8 file names is now available for Master Files, Data Files, and Application Directory names.

Unicode PDF Output

Reporting from a Unicode data source with PDF output format is available when using the following two fonts which support Unicode characters:

- ❑ Lucida Sans Unicode, which is used to display single-byte characters only. This font is available on Windows version 2000 and higher.
- ❑ Arial Unicode MS, which is used to display both single-byte and double-byte characters. This font can be installed as an option from the Microsoft Office CD, version 2000 and higher.

The Lucida Sans Unicode font is the default font if the WebFOCUS Reporting Server is configured for UTF-8 (code page 65001) or UTF-EBCDIC (code page 65002).

If the server is configured for Unicode, and you want to use the Arial Unicode MS font instead, you must specify the Arial Unicode MS font in the StyleSheet. Alternatively, if the WebFOCUS PDF font mapping file (for example, EDAHOME NLS pdf.fmp) has DEFAULT-FONT=YES specified for Arial Unicode MS, then it becomes the default font. For example:

```
font=Arial Unicode MS, style=normal, metricsfile=PDARUM AFM *,  
DEFAULT-FONT=YES, $
```

For more information about how to use the PDF font mapping file, see the *Creating Reports With TIBCO WebFOCUS® Language* manual.

EDATEMP and NFS-Mounted Disks

This appendix describes the Reporting Server objects in EDATEMP for Interprocess Communication (IPC).

In this appendix:

- [EDATEMP and NFS-Mounted Disks](#)
-

EDATEMP and NFS-Mounted Disks

Part of the Reporting Server process creates pipe and fifo objects in EDATEMP for Interprocess Communication (IPC). However, pipe and fifo object types for IPC use do not work well on NFS-mounted disks. To alleviate this problem and allow NFS-mounted disk use, the EDAIPC system variable may be set to a local non-NFS directory location, so pipe and fifo object creations are redirected to this separate non-NFS location.

Normally, the EDATEMP directory defaults to a subdirectory of EDACONF. The EDATEMP variable allows EDACONF file objects, such as edaprint.log and rmdta.log and the EDATEMP directory to be redirected to the location specified by EDATEMP. EDAIPC is usually used in conjunction with EDATEMP, when EDATEMP is pointed to an NFS-mounted disk. It is also possible that EDACONF itself could be on an NFS-mounted disk (the EDATEMP directory is also on an NFS-mounted disk by parentage) and EDAIPC should also be used in this instance.

For example:

```
export EDATEMP=/nfs/wfs
export EDAIPC=/tmp/wfs
.../edastart -start
```

The /tmp directory is always local, but should organizationally be identified by a subdirectory name, such as wfs. If multiple Reporting Servers are run on the same box, there should also be an identify subdirectory such as wfs/accounting to prevent process collisions.

Note: The pipes and fifos for tscom300 are the only exceptions to the redirection that EDAIPC invokes. Additionally, while the Reporting Server can detect if NFS is in use, EDAIPC has no default and must be explicitly set to redirect the pipes and fifos.

The EDATEMP and EDAIPC variables may also be set in the EDACONF bin/edaenv.cfg file, using the normal edaenv.cfg rule, which is one variable per line in x=y format (no "export" keyword).

Caching Support for Node.js

Node Package Manager (npm) is a package manager for the Reporting Server Node.js[®] caching server run-time environment. Node Package Manager and Node.js are third-party software packages that should, preferably, be installed before you install the Reporting Server so that automatic configuration and usage occurs.

On Windows, npm and Node.js are packaged together. Other platforms typically require npm to be installed first and then for npm to install Node.js second.

Note: Throughout this topic, the npm software and Node.js caching server environment are referred to as Node.js.

In this appendix:

- [Node.js Installation and Configuration](#)
 - [Node.js Process](#)
 - [Node.js Prerequisites](#)
-

Node.js Installation and Configuration

When Node.js is installed and running, and a Reporting Server is configured to use it, In-Document Analytics reports are cached under Node.js and can be accessed independent of the Reporting Server.

If Node.js is installed on Windows or Linux/UNIX and is detected as available during Reporting Server installation time, the In-Document Analytics Caching feature will automatically be configured in the EDACONF etc/ar_v2 directory and registered in the edaserve.cfg file, using a port in the reserved port range for the Reporting Server (HTTP port number plus 5).

The In-Document Analytics Caching feature can also be configured post-installation, with a few simple steps as outlined below, if:

- Node.js was not detected during installation, due to not being installed or not added to the PATH environment variable.
- Node.js configuration failed during automatic configuration.
- Reporting Server binaries were upgraded to a level that supports Node.js caching and the earlier Reporting Server did not support Node.js caching or was not configured.

For performance reasons, it is always preferable to install on a local disk. However, as of Release 8207.28, it is possible to use NFS if the EDAIPC environment variable is used during server run time to point at a local disk location for pipe and fifo creation. For more information on setting EDAIPC, see [EDATEMP and NFS-Mounted Disks](#) on page 833.

Node.js Process

Starting the Reporting Server checks if the configured Node.js port is already available, and therefore Node.js is already running. If the Node.js process is not running, the Reporting Server will start the Node.js caching server process. The Node.js caching server process has the name `node` or `node.exe`, depending on your platform. If the Reporting Server starts the Node.js process, it will also stop the process on server shutdown. In this way, you may control your own Node.js usage or let the Reporting Server be in control.

By default, the Reporting Server process is configured to point to a URL for a local Node.js In-Document Analytics Caching process. As such, the Node.js URL may be reconfigured to run against a remote machine, provided the remote machine is properly configured and running. If a remote URL is used, you are expected to start the remote Node.js server process before the local Reporting Server, as the local server cannot start a remote Node.js process. The simplest form of remote configuration is to point to another Reporting Server Node.js In-Document Analytics Caching URL and let that server's start up control the Node.js start up. A more complex variation is to create a *freestanding* Node.js In-Document Analytics Caching environment with no Reporting Server involved and then point any number of Reporting Servers to it. For more information, see [How to Manually Configure a Freestanding Node.js In-Document Analytics Caching Environment on a Separate Machine](#) on page 841.

Note: On Linux/UNIX, if the installation ID has a non-writable root owned \$HOME/.npm directory tree, the automatic configuration, as well as manual configuration, will fail for both the internal `npm install` step and the Reporting Server's Node.js startup. The solution for this is to delete the \$HOME/.npm directory tree or use an installation ID that does not have a \$HOME/.npm directory.

Node.js Prerequisites

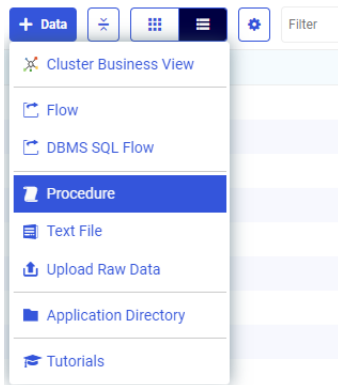
The following are the prerequisites for Node.js:

- ❑ **For Windows.** Download and run the Node.js Windows 64-bit installer from <https://nodejs.org/en/download>. You can use any recent release. Take the default installation options. Do not deselect the *Add to PATH* option and do *not* check *Automatically install the necessary tools* option, unless you have need for this option for other software that may also be using this Node.js instance.

- ❑ **Linux/UNIX.** Most Linux/UNIX vendors will have their own npm and Node.js installer. Any recent release may be used. The installation steps may be found by performing an internet search on *install nodejs on {distribution name such as RedHat or Ubuntu}*. The npm portion is typically a separate step, but noted within the Node.js package instructions.

Procedure: How to Verify the Node.js In-Document Analytics Caching Environment is Running

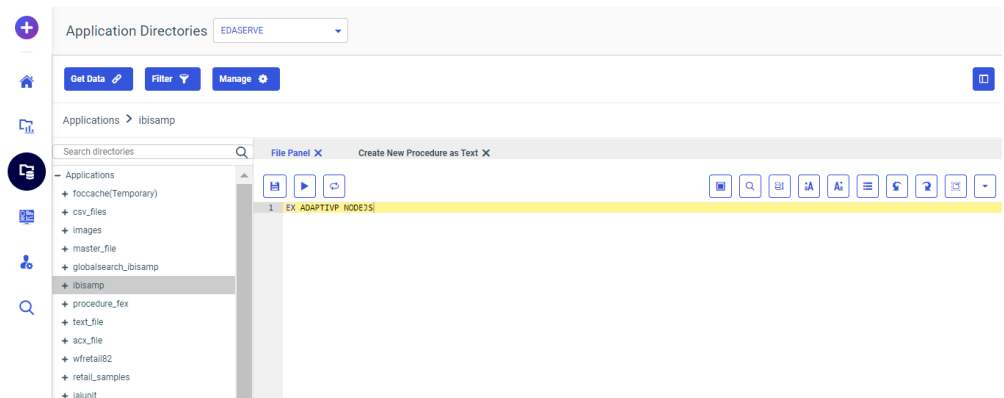
1. Start the Reporting Server, if it is not already running.
2. Sign in to the Reporting Server browser interface using any ID that has open (create) procedure authority.
3. Click + *Data* and select *Procedure* from the menu to create a new procedure, as shown in the following image.



4. In the procedure text window, type the following command:

```
EX ADAPTIVP NODEJS
```

An example of the procedure text window is shown in the following image.



- Click the run button ▶.

Note: There is no need to save the procedure first.

An In-Document Analytics-based report should appear, as shown in the following image.

1 - 57 of 157 records Page: 1 / 3

AR NODEJS Caching Manager Test								
Table of supported Currencies								
Currency	ISO code	Symbol	Symbol hexadecimal	Symbol Char-length	Symbol Byte-length	Decimals	Position in report	Non-UTF Symbol
USA dollar	USD	\$	24	1	1	2	4	
Eurozone euro	EUR	€	80	1	1	2	3	
Japanese yen	JPY	¥	a5	1	1	0	2	
British pound sterling	GBP	£	a3	1	1	2	1	
United Arab Emirates dirham	AED			0	0	2	0	
Afghan afghani	AFN	Afs	41, 66, 73	3	3	2	0	Afs
Albanian lek	ALL	Lek	4c, 65, 6b	3	3	2	0	
Armenian dram	AMD			0	0	2	0	
Netherlands Antillean guilder	ANG	f	83	1	1	2	0	f

If the Node.js process is not working, a *refuse to connect* message containing the URL host name will appear.

If the test fails, check for *nodejs startup failed* in the edaprint file, during initial startup of the Reporting Server. This typically happens if Node.js was not installed, Node.js is not on the PATH, or the initial automatic or manual configuration failed. Follow the steps in [How to Add Node.js Configuration to a Reporting Server Not Originally Configured With Node.js](#) on page 839 to investigate and correct.


Another possible failure is the Node.js process is no longer running. To check:

- On Linux/UNIX, issue the following command to limit output:

```
ps -edf | grep "node"
```

Look for processes, such as:

```
webfocus 122264 122253 0 10:20 ? 00:00:00 sh -c set DEBUG=http,express:*
& node --max-old-space-size=16192 server.js "--PORT=8126"
webfocus 122266 122264 0 10:20 ? 00:00:00 node --max-old-space-
size=16192 server.js --PORT=8126
```

- On Windows Task Manager, find the process icon  Node.js JavaScript Runtime and turn on the Task Manager command line column, which will show something similar to the following.

```
node --max-old-space-size=16192 server.js --PORT=8126
"C:\Program Files\nodejs\node.exe" "C:\Program Files\nodejs\node_modules\npm\bin\npm-cli.js" run mystart --prefix=c:\ibi\srv90\wfs\etc\ar_v2\cachemanager -- --PORT=8126
```

- ❑ Check for issues, such as a wrong port number or the process is not running. If not running, check the EDACONF edatemp/npmrun.trc file for the possible reason.

Procedure: How to Add Node.js Configuration to a Reporting Server Not Originally Configured With Node.js

Note: Be prepared to recycle the Reporting Server as part of this procedure. You may want to do this during scheduled maintenance down time.

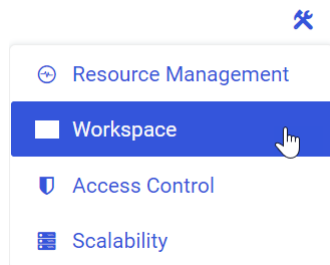
1. Sign in to the Reporting Server browser interface, using the server administrator ID.
2. Click + *Data* and select *Procedure* from the menu to create a new procedure.
3. In the procedure text window, type the following command:

```
EX ADAPTIVE NODEJS,REBUILD
```

4. Click the run button ▶.

Note: There is no need to save the procedure first. It is normal to see warnings and npm update failed messages. Any other fail messages should be investigated and reported.

5. Navigate to the Workspace page by clicking the *Tools* icon in the banner and selecting *Workspace*, as shown in the following image.



6. Under Data Services on the tree, open the *Configuration Files* folder, and double-click *Workspace - edaserve.cfg* to edit the file.
7. In the [Workspace] section add, if not already present, or correct the following nodejs_url line:

```
nodejs_url = http://{hostname}:{port}
```


where:

hostname

May be localhost or a specific host name, with or without the domain portion. The host name does not need to be the same machine as the Reporting Server, but the automatic Node.js start feature will not work if it is not on the same machine. If a remote Node.js configuration is in use, it is expected to be running before the Reporting Server is started.

port

Is typically the Reporting Server HTTP port number plus 5, if automatic configuration was successful. For example, HTTP port number 8101 would be 8106 for Node.js, but it may be any free port. The Node.js default port number is 8090, if you wish to skip the "--PORT=####" parameter in the manual start instructions.

8. Click the Save button  , which will save and restart the Reporting Server. If the nodejs_url line was already present and did not need edits, still click the Save button to force a restart of the Reporting Server.

Procedure: How to Manually Start and Stop Node.js In-Document Analytics Caching

To manually start Node.js In-Document Analytics Caching, you will need to know:

- Where the *npm* command is installed on your machine if it is not on \$PATH (Linux/UNIX) or %PATH% (Windows) and add that path to the npm commands below (for example, /usr/local/bin/npm).
- Full path directory for the server EDACONF etc/ar_v2/cachemanager directory.
- A writable directory for output. This is normally in the EDACONF edatemp directory when a Reporting Server starts Node.js. This directory is not advisable in a privately started environment. From an organizational perspective, EDACONF would be a reasonable location.

The following are examples of starting Node.js In-Document Analytics Caching. Replace the applicable paths and port number with your paths and port number.

Note: These commands must be typed on a single line.

For Linux/UNIX:

```
npm run mystart --prefix=/webfocus/ibi/srv90/wfs/etc/ar_v2/cachemanager --  
--PORT=8126 > /webfocus/ibi/srv90/wfs/npmrun.trc 2>&1 &
```

For Windows:

```
npm run mystart --prefix=c:\ibi\srv90\wfs\etc\ar_v2\cachemanager -- --
PORT=8126 > c:\ibi\srv90\wfs\npmrun.trc 2>&1
```

On Linux/UNIX, a `ps -edf` command will show something similar to the following:

```
webfocus 122264 122253 0 10:20 ? 00:00:00 sh -c set DEBUG=http,express:* &
node --max-old-space-size=16192 server.js "--PORT=8126"
webfocus 122266 122264 0 10:20 ? 00:00:00 node --max-old-space-size=16192
server.js --PORT=8126
```

On Windows Task Manager, find the  Node.js JavaScript Runtime process icon and turn on the Task Manager command line column, which will show something similar to the following.

```
node --max-old-space-size=16192 server.js --PORT=8126
"C:\Program Files\nodejs\node.exe" "C:\Program Files\nodejs\node_modules
\npm\bin\npm-cli.js" run mystart --prefix=c:\ibi\srv90\wfs\etc
\ar_v2\cachemanager -- --PORT=8126
```

The following are examples for stopping Node.js In-Document Analytics Caching. Replace the applicable paths and port number with your paths and port number.

Note: These commands must be typed on a single line.

For Linux/UNIX:

```
npm stop mystart --prefix=/webfocus/ibi/srv90/wfs/etc/ar_v2/cachemanager > /
webfocus/ibi/srv90/wfs/npmrun.trc 2>&1 &
```

For Windows:

```
npm stop mystart --prefix=c:\ibi\srv90\wfs\etc\ar_v2\cachemanager > c:\ibi
\srv90\wfs\npmrun.trc 2>&1
```

Procedure: How to Manually Configure a Freestanding Node.js In-Document Analytics Caching Environment on a Separate Machine

Note: This assumes Node.js is already installed and on the PATH of the separate machine.

1. Navigate to the EDASHOME etc/etc/ar_v2 directory on a Reporting Server where Node.js In-Document Analytics Caching is supported.
2. Archive, using zip or tar, the following files for transfer to the target new machine:
 - ar_v2/cachemanager/*.js
 - ar_v2/cachemanager/package.json
 - ar_v2/cachemanager/template/stored.html

3. Restore the files to a writable directory on the new machine, for example, on Linux/UNIX:

```
$HOME/webfocus/ar_v2
```

4. Edit the `ar_v2/cachemanager/config.json` file.

- ❑ On the `redirect_host` line, change the `redirect_host http` value, usually `localhost`, to the actual full DNS machine name, for example:

```
"redirect_host": "http://mymachine.mycompany.com"
```

- ❑ On the `cache_port` line, change the `cache_port` number to the port number you want to use for the configuration, matching the port used by the Reporting Server `nodejs_url` value.

5. Change the directory to the `cachemanager` directory, for example, on Linux/UNIX:

```
cd $HOME/webfocus/ar_v2/cachemanager
```

6. Run the following command:

```
npm install
```

7. Review the installation for any errors. It is normal to see warnings and `npm update` failed messages. Any other fail messages should be investigated and reported.

Note: A root owned `$HOME/.npmrc` directory is typically a reason for failure and should be deleted beforehand or, another ID, without a root owned `$HOME/.npmrc` directory, should be used.

8. Start and stop the environment, as described in [How to Manually Start and Stop Node.js In-Document Analytics Caching](#) on page 840, using the applicable paths and ports.

Procedure: How to Block Automatic Node.js Startup By the Reporting Server

If you want to stop the server from attempting automatic Node.js startup, even though it is configured, you can set the `BLOCK_NODEJS` environment variable to a value of `Y`.

The best place to do this is in the `EDACONF bin/edaenv.cfg` file by adding the `"BLOCK_NODEJS=Y"` line, and saving and restarting the Reporting Server. It may also be set as an environment variable, and exported on Linux/UNIX, before Reporting Server startup, although this is less convenient for Windows Reporting Servers running as a service, as it needs to be set up at a system level.

Note: The `edaenv.cfg` file may not exist. If this is the case, create the file and add the `"BLOCK_NODEJS=Y"` line.

Platform-Specific Methods for Specifying EDAEXTSEC

Environment variable EDAEXTSEC is only used to override the security provider setting that is stored in edaserve.cfg file in order to start the server with security OFF.

In this appendix:

- [Platform-Specific Methods for Setting EDAEXTSEC](#)

Platform-Specific Methods for Setting EDAEXTSEC

Platform	Start Method	Actions
Windows	Start Menu	Under the main software folder for the server, select <i>Start Security ON</i> or <i>Start Security OFF</i> .
	As a Service under a local system account	Set the system environment variable EDAEXTSEC, under <i>My Computer, Properties, Advanced</i> tab, to OFF. Restart Windows to initialize.
	As a Service under the current account	Set the system environment variable EDAEXTSEC, under <i>My Computer, Properties, Advanced</i> tab, to OFF. Restart Windows to initialize.
UNIX	Command prompt or Script that calls edastart	Export the variable setting in the edastart shell script, or before any calls to the shell script, as follows: <code>export EDAEXTSEC=OFF</code>
z/OS Unified Server	JCL	Place the variable in the EDAENV allocation of the IRUNJCL member of the configuration library <code>EDAEXTSEC=OFF</code>

Platform	Start Method	Actions
IBM i	QSH command prompt or QSH script that calls edastart	Export the variable in the edastart shell script, or before any calls to the shell script, as follows: <code>export EDAEXTSEC=OFF</code>
	CALL {lib}/TSCOM300 or CL that calls TSCOM300 or CL that calls QSH, which, in turn, calls the edastart shell script	Before the call to TSCOM300, add the following: <code>ADDENVVAR ENVVAR(EDAEXTSEC) VALUE(OFF)</code>
Any	Any	Use the Reporting Server browser interface (<i>Workspace</i> , <i>Configuration Files</i> folder, <i>Miscellaneous</i> folder, <i>Environment - edaenv.cfg</i>) to edit the edaenv.cfg file and add the following: <code>EDAEXTSEC=OFF</code> Note: After a change, a hard restart of the server must be done.

Legal and Third-Party Notices

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, the TIBCO logo, the TIBCO O logo, FOCUS, iWay, Omni-Gen, Omni-HealthData, and WebFOCUS are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

This software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the readme file for the availability of this software version on a specific operating system platform.

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (<https://www.tibco.com/patents>) for details.

Copyright © 2022. TIBCO Software Inc. All Rights Reserved.

Index

- ? command [496](#)
- * command [496](#)
- CLOSE command [496](#)
- DEFAULTS command [497](#), [511](#)
- EXIT command [497](#)
- GOTO command [497](#), [513](#)
- IF command [497](#), [514–516](#)
- INCLUDE command [497](#), [516](#), [517](#)
- label command [497](#)
- PASS command [497](#)
- PROMPT command [497](#)
- QUIT command [497](#)
- READ command [497](#)
- REMOTE BEGIN command [497](#)
- REMOTE END command [497](#)
- REPEAT command [498](#)
- RUN command [498](#)
- SET command [498](#), [512](#)
- TYPE command [498](#)
- WRITE command [498](#)
- &&name variable [502](#)
- &APPROOT variable [504](#)
- &DATE variable [504](#)
- &DATEfmt variable [504](#)
- &DMY variable [505](#)
- &DMYY variable [505](#)
- &ECHO variable [505](#), [511](#)
- &EXITRC variable [506](#)
- &FOCFOCEXEC variable [505](#)

- &FOCINCLUDE variable [505](#)
- &FOCMODE variable [505](#)
- &FOCPRINT variable [505](#)
- &FOCREL variable [506](#)
- &IORETURN variable [506](#)
- &MDY variable [504](#)
- &MDYY variable [505](#)
- &name variable [502](#)
- &RETCODE variable [506](#)
- &TOD variable [507](#)
- &YMD variable [505](#)
- &YYMD variable [505](#)

A

- accessing APP directories on Windows [182](#)
- accessing pop-up menus from Reporting Server browser interface navigation panes [20](#)
- active connections [674](#)
- adapter configuration [236](#)
 - changing [236](#)
- adapter connections [236](#)
- adapters [234](#), [272](#)
 - complete list report [238](#)
 - configuring [234](#)
 - data type report [242](#)
 - license counts [243](#)
 - SQL APT status report [239](#)
- administering servers [17](#), [18](#), [573](#)
- administration levels [80](#)

- administrative tasks for Data Services [676](#)
- administrative tasks for Java Services [687](#)
- administrative tasks for Special Services and Listeners [700](#)
- agent monitoring [660](#), [668](#), [684](#), [690](#)
- agent services [576](#)
 - configuring [576](#)
- agents [663](#), [678](#), [680](#)
 - error messages [680](#)
 - starting [680](#), [688](#)
 - statistics [663](#)
 - stopping [678](#), [679](#)
- aggregated statistics for Data Services [669](#)
- aggregated statistics for Java Services [686](#)
- ALLOCATE subcommand [531](#)
- allocating data files [214](#)
- allocating data sets [531](#)
- allocating files [549](#)
- allocating temporary files [220](#)
- alphanumeric expressions [521](#), [522](#), [524](#)
- AND keyword [523](#)
- answer sets
 - returning [518](#), [519](#), [526](#), [527](#)
- APP ? METALLOCATION_SAME command [183](#)
- APP APPENDPATH command [179](#)
- APP application management commands [175](#)
- APP commands [174](#)
 - APP ? METALLOCATION_SAME [183](#)
 - APP APPENDPATH [177](#), [179](#)
 - APP COPY [186](#)
 - APP commands [174](#)
 - APP COPYF [187](#)
 - APP CREATE [175](#)
 - APP DELETE [175](#), [190](#)
 - APP DELETEDF [190](#)
 - APP FI [201](#)
 - APP HELP [209](#)
 - APP HOLD [199](#)
 - APP LIST [203](#)
 - APP MAP [180](#)
 - APP MOVE [188](#)
 - APP MOVEF [188](#)
 - APP PATH [177](#), [178](#)
 - APP PREPENDPATH [177](#), [178](#)
 - APP PROPERTY CODEPAGE [191](#)
 - APP QUERY [204](#)
 - APP RENAME [175](#), [191](#)
 - APP RENAMDF [192](#)
 - APP SET METALLOCATION_SAME [182](#)
 - APP SHOWPATH [183](#)
 - FI [197](#)
 - file types [193](#)
 - HOLD [197](#)
 - HOLDDATA [197](#), [200](#)
 - HOLDMETA [197](#), [200](#)
 - APP COPY command [186](#)
 - APP COPYF command [187](#), [193](#)
 - APP DELETE command [190](#)
 - APP DELETEDF command [190](#), [193](#)
 - APP FI command [197](#), [201](#)

- APP HELP command [176, 209](#)
 - APP HOLD command [197, 199](#)
 - APP HOLDDATA command [197, 200](#)
 - APP HOLDMETA command [197, 200](#)
 - APP LIST command [176, 203](#)
 - APP MAP command [180](#)
 - APP MOVE command [188](#)
 - APP MOVEF command [188, 193](#)
 - APP PATH command [160, 178](#)
 - configuring [178](#)
 - search paths [178](#)
 - APP PREPENDPATH command [178](#)
 - APP PROPERTY CODEPAGE command [191](#)
 - APP QUERY command [176, 204](#)
 - APP RENAME command [191](#)
 - APP RENAMEF command [192, 193](#)
 - APP reporting and help commands [176](#)
 - APP SET METALLOCATION_SAME command [182](#)
 - APP SHOWPATH command [183](#)
 - appearance of Reporting Server browser interface [25](#)
 - Appearance option [25](#)
 - appearance, web console [590](#)
 - appending applications to search path [179](#)
 - application (APP) [118](#)
 - physical location of [118](#)
 - application components [117, 118](#)
 - listing [204](#)
 - application directories [134, 160, 163, 180, 184](#)
 - creating [184](#)
 - application directories [134, 160, 163, 180, 184](#)
 - filtering in Reporting Server browser interface [163](#)
 - mapping from Reporting Server browser interface [134](#)
 - mapping manually [180](#)
 - application paths [177](#)
 - APP PATH [160](#)
 - APP PATH command [178](#)
 - configuring [177](#)
 - application repositories [117, 118](#)
 - application settings [620](#)
 - applications [190](#)
 - deleting [190](#)
 - APPROOT hierarchy [75](#)
 - aproot parameter [118](#)
 - arithmetic expressions [520, 521](#)
 - averages [673](#)
- B**
- baseapp directory [119](#)
 - binning in data assist [326](#)
 - binocular search [457](#)
 - branching [513, 514](#)
- C**
- caching files for WebFOCUS Active Technologies [620](#)
 - calling procedures [516, 517](#)
 - case translation functions [828](#)

- changing adapter configurations [236](#)
 - changing passwords [61](#)
 - changing the password delimiter [60](#)
 - character semantics and functions [828](#)
 - CLOSE subcommand [542](#)
 - cluster joins [273](#)
 - Cluster Manager [707](#)
 - Cluster Manager properties [718](#)
 - code page values [598](#)
 - for IBM i [600](#)
 - for UNIX [598](#)
 - for Windows [598](#)
 - codepage, application [191](#)
 - command lines [511](#)
 - commands [495](#)
 - communications configuration [238](#)
 - compound expressions [525](#)
 - COMPRESS subcommand [549](#)
 - compute in synonym editor [299](#)
 - concatenating files [550](#)
 - concatenation [522](#)
 - conditional branching [514](#)
 - configuration files [118](#)
 - configuration files, editing [701](#)
 - configuring [576](#)
 - agent services [576](#)
 - deployment modes [576](#)
 - Connect to Data [267](#)
 - connection monitoring [674](#), [675](#)
 - Connection Queue [674](#), [675](#)
 - connection statistics [675](#)
 - connections [675](#)
 - monitoring [675](#)
 - stopping [675](#), [676](#)
 - CONTAINS operator [523](#)
 - controlling HLIPRINT output [767](#)
 - controlling output [767](#)
 - cookie, retrieving the value of [106](#)
 - COPY subcommand [543](#)
 - COPYDD subcommand [545](#)
 - core engine global variables [609](#)
 - core engine settings [605](#)
 - Core Engine settings [605](#)
 - core system variables [608](#)
 - creating data files [214](#)
 - creating data sets [550](#)
 - creating expressions [519](#)
 - creating procedure files [212](#)
 - creating synonyms [245](#)
 - CURRENCY_PRINT_ISO parameter [596](#)
 - custom message from quiesced server [575](#)
 - customer support using trace files [794](#)
- D**
- data access agents [576](#)
 - data assist [270](#), [275](#), [321](#)
 - binning [326](#)
 - business view frame [285](#)
 - common options [279](#)
 - data profiling [331](#)

- data assist [270](#), [275](#), [321](#)
 - DBA [314](#)
 - decode [292](#)
 - expression calculator [301](#)
 - filters [321](#)
 - filters/groups frame [287](#)
 - geographic role [328](#)
 - grouping [323](#)
 - joins [308](#)
 - output pane [293](#)
 - panes [279](#)
 - pivoting columns [319](#)
 - ribbon [275](#)
 - sample data [330](#)
 - table/column pane [280](#)
 - variables [306](#)
- data files [214](#)
- data flows [345](#)
 - creating [345](#)
 - preparing data [376](#)
- data preparation [376](#)
- data profiling in data assist [331](#)
- data set names [219](#)
- data sets [528](#)
 - allocating [531](#)
 - creating [550](#)
 - locking [530](#)
- data, accessing [271](#)
- Data
 - preparing [257](#)
- DATE_ORDER parameter [596](#)
- DATE_SEPARATOR parameter [597](#)
- DBA in data assist [314](#)
- DBMS security [56](#)
- ddnames (logical names) [558](#)
- decoding, comparison and editing functions [828](#)
- deferred files [743](#)
 - valid extensions [743](#)
- Deferred List [695](#)
- deferred list statistics [696](#)
- deferred management [735](#)
 - setting properties in [735](#)
- deferred procedure execution [731](#), [735](#), [742](#)
- deferred process [742](#)
 - monitoring [742](#)
 - statistics [742](#)
- deferred requests [695](#), [696](#)
 - deleting [697](#)
 - monitoring [695](#), [696](#)
 - retrieving answer sets [697](#)
- define in synonym editor [299](#)
- DELETE subcommand [546](#)
- deployment modes [580](#)
- designating a network drive on Windows [182](#)
- diagnostics for servers [757](#), [792](#)
 - EDAPRINT files [759](#)
 - record and playback [781](#), [787](#), [791](#)
 - trace files [768](#)
- Dialogue Manager commands [496](#), [528](#)
 - ? command [496](#)

Dialogue Manager commands [496](#), [528](#)

- * command [496](#)
- CLOSE [496](#)
- DEFAULTS [497](#), [511](#)
- EXIT [497](#)
- GOTO [497](#), [513](#)
- IF [497](#), [514–516](#)
- INCLUDE [497](#), [516](#), [517](#)
- label [497](#)
- PASS [497](#)
- PROMPT [497](#)
- QUIT [497](#)
- READ [497](#)
- REMOTE BEGIN [497](#)
- REMOTE END [497](#)
- REPEAT [498](#)
- RUN [498](#)
- SET [498](#), [512](#)
- TYPE [498](#)
- WRITE [498](#)

Dialogue Manager procedures [495](#)

- calling [499](#)
- platform-specific commands [528](#)
- running [498](#)
- variables and [501](#), [502](#)

dimension [270](#)

DYNAM command [215](#), [528](#), [530](#), [549](#)

DYNAM SET LONGSNYM [222](#)

DYNAM subcommands [528](#)

- ALLOCATE [531](#)

DYNAM subcommands [528](#)

- CLOSE [542](#)
- COMPRESS [549](#)
- COPY [543](#)
- COPYDD [545](#)
- DELETE [546](#)
- RENAME [547](#)
- SUBMIT [548](#)

DYNAM SUBMIT command [625](#)

E

EDACONF [74](#), [75](#)

EDAHLIP variable [767](#)

EDAHOME [74](#), [75](#)

EDAPRINT files [759](#)

EDASERVE.CFG [234](#)

edaserve.cfg configuration file [118](#)

EDASPROF.PRF [234](#)

EDATEMP [74](#), [75](#)

editing [701](#)

- server files [701](#)

- user.prf file [701](#)

email SMTP server settings [617](#)

enabling Cluster Manager [707](#)

encrypting passwords [97](#)

encrypting passwords for configuration files [97](#)

EQ operator [523](#)

error messages [680](#)

- agents [680](#)

event routing [582](#)

EX procedures [226](#)

EXEC command [509](#), [516](#), [517](#)

execution flow [511](#), [513](#)

- debugging [511](#)

expression calculator in synonym editor [301](#)

expressions [519–525](#)

F

FDS (FOCUS Database Server) [766](#)

file attributes

- specifying [558](#)

FILEDEF command [214](#), [215](#), [219](#), [558](#), [560](#), [561](#)

FILEDEF

- CONCAT [215](#)

files [549](#)

- concatenating [550](#)
- freeing [550](#)

filtering application directories in Reporting Server

- browser interface [163](#)
- filtering traces [772](#)

filters in data assist [321](#)

fine-tuning the server [704](#)

- JVM memory usage [705](#)

Fixed files [827](#)

Fixed files configuration for Unicode [827](#)

flow of execution [511](#), [513](#)

- debugging [511](#)

flows [345](#)

foccache directory [119](#), [620](#)

FOCUS Database Server (FDS) [766](#)

focus limits, displaying [611](#)

focus sets and info [605](#)

fonts [831](#)

- Unicode in PDFs [831](#)

Force Finalization [706](#)

format conversion functions [828](#)

freeing files [550](#)

function parameters, displaying [615](#)

G

GE operator [523](#)

General Workspace Configuration Settings [625](#)

- calling exit for DYNAM SUBMIT on z/OS [625](#)

geographic role in data assist [328](#)

GETCOOKI function [106](#)

global server profiles [28](#)

global variables [501](#), [503](#), [609](#)

- naming [501](#), [502](#)

group profiles [31](#)

grouping in data assist [323](#)

GT operator [523](#)

H

header variable, retrieving the value of [106](#)

heap size [705](#)

HLIPRINT log [766](#), [767](#)

- controlling [766](#)
- controlling number of lines [766](#)

HOLD keywords, displaying [610](#)

HTI script (HTTP internal script) [787](#)

hto files [791](#)

HTTP header variable, retrieving the value of [106](#)

HTTP internal script (HTI script) [787](#)

I

ibisamp directory [119](#)

IBISNAP [759](#)

IBM i language and code page values [600](#)

J

Java Services [681](#), [682](#)

java services, configuring [702](#)

Java Virtual Machine (JVM) [706](#)

JCL commands [549](#)

joins in data assist [308](#)

JSCOM3 connections [681](#)

justification functions [828](#)

JVM [705](#)

JVM (Java Virtual Machine) [706](#)

memory usage [705](#)

tuning options [705](#)

K

keyword parameters

passing [509](#)

killing agents [678](#)

by user id [679](#)

L

language values [598](#)

for IBM i [600](#)

for Windows and UNIX [598](#)

languages and code pages [598](#)

for IBM i [600](#)

for Windows and UNIX [598](#)

LDAP security [56](#)

LDAP security settings [56](#)

LDAP support [56](#)

LE operator [523](#)

left-click menus [20](#)

length and position functions [828](#)

limits, displaying [611](#)

listener statistics [694](#)

listeners [689](#)

Listeners and Special Services Configuration page
[727](#)

local variables [501](#), [502](#)

naming [501](#), [502](#)

locale settings [591](#)

locking data sets [530](#)

logical expressions [523](#), [524](#)

logical names [558](#)

logical operators [523](#)

Login Info option [24](#)

login information [24](#)

login message [587](#)

login password expiration [61](#)

LONGSYNM [222](#)

M

managing deferred requests [735](#), [742](#)
managing servers [17](#), [18](#), [573](#)
mapping application directories [134](#)
 manually [180](#)
Master Files [219](#)
max_connections_per_user [579](#)
max_connections_per_user property [579](#)
measure [270](#)
messages, displaying [613](#)
metadata [211](#), [270](#)
 accessing [211](#)
Microsoft SQL Server configuration for Unicode
[823](#)
migrate settings [658](#)
migrating servers [658](#)
miscellaneous settings [622](#)
modes of server deployment [580](#)
monitoring [742](#)
 deferred [742](#)
MVS commands [528](#)
My Console options [23](#)

N

naming variables [501](#), [502](#)
NE operator [523](#)
nesting [517](#)
non-data files [212](#)
NOT operator [523](#)
notifying about password expiration [61](#)

O

OFFLINE command [561](#)
OMITS operator [523](#)
ON TABLE HOLD [518](#), [526](#)
ON TABLE PCHOLD [518](#), [519](#), [526](#), [527](#)
opening the Reporting Server browser interface [18](#)
operators [523](#)
 CONTAINS [523](#)
 EQ [523](#)
 GE [523](#)
 LE [523](#)
 NOT [523](#)
 OMIT [523](#)
 OR [523](#)
OPSYS security [46](#)
 APPROOT [75](#)
 EDACONF [75](#)
 EDAHOME [75](#)
 EDATEMP [75](#)
 file areas [74](#)
 temporary files (EDATEMP) [75](#)
OR operator [523](#)
Oracle character semantics [821](#)
Oracle configuration for Unicode [821](#)
output files [199](#), [200](#)
 creating [199](#), [200](#)

P

parameters
 passing [509](#), [510](#)

- passing parameters [509](#), [510](#)
 - password delimiter [60](#)
 - password encryption for configuration files [97](#)
 - password expiration notification [61](#)
 - passwords
 - changing [61](#)
 - encrypting [97](#)
 - PDF report output [831](#)
 - Unicode fonts [831](#)
 - Performance Report for Data Service Agents [661](#), [663](#)
 - Performance Report for Java Service Agents [681](#)
 - Performance Report for Special Services and Listeners [689](#), [690](#)
 - Performance Report for Web Sessions [697](#), [698](#)
 - physical file location [201](#)
 - mapping to [180](#)
 - pivot in data assist [319](#)
 - platform-specific commands [528](#)
 - playback hto files [791](#)
 - playback log files [791](#)
 - playback of user actions [787](#), [791](#)
 - [781](#)
 - Playback tool features [786](#)
 - percentile processing in stress tests [786](#)
 - refreshing directories [786](#)
 - pooled server deployment mode [580](#)
 - positional parameters [510](#)
 - passing [510](#)
 - preference settings for Reporting Server browser interface [23](#)
 - Preferences option [25](#)
 - prepending applications to search path [178](#)
 - private server deployment mode [580](#)
 - procedures [212](#), [455](#)
 - calling [516](#), [517](#)
 - creating [456](#)
 - editing [459](#)
 - printing [459](#)
 - run-time processing [495](#)
 - running [471](#)
 - scheduling [488](#)
 - testing [511](#)
 - profile levels [28](#)
 - profile settings [618](#)
 - profiles [177](#)
 - applications and [177](#)
 - PTH security [51](#)
- Q**
- queue_max_user_conns [579](#)
 - queue_max_user_conns property [579](#)
 - queued connections [674](#)
 - Queuing property [579](#)
 - queuing user connections [579](#)
 - quiesced server [575](#)
- R**
- recording user actions [781](#), [787](#), [791](#)

- releasing unused memory [706](#)
 - remote servers [238](#), [711](#)
 - remote servers, adding [711](#)
 - RENAME subcommand [547](#)
 - Reporting Server browser interface [17](#), [18](#), [178](#), [573](#)
 - application paths [178](#)
 - Metadata tree [163](#)
 - troubleshooting [757](#), [759](#), [768](#), [781](#), [787](#), [791](#), [792](#)
 - repository procedure, running [473](#)
 - right-click menus [20](#)
 - role profiles [31](#)
 - Run Options [25](#)
 - running procedures [455](#)
 - running the WebFOCUS Reporting Server [573](#)
- S**
- sample data in data assist [330](#)
 - SAP BW configuration for Unicode [826](#)
 - SAP R/3-ECC configuration for Unicode [826](#)
 - Save Last Report option [24](#)
 - savediag facility [768](#), [794](#)
 - saving last report [24](#)
 - Scheduler agents [731](#)
 - script file [773](#)
 - capturing traces [773](#)
 - search paths [178](#)
 - adding names to [178](#), [179](#)
 - appending applications manually [179](#)
 - search paths [178](#)
 - commands [174](#)
 - prepending applications manually [178](#)
 - search, binocular [457](#)
 - security DBMS [56](#)
 - security features [80](#), [87](#)
 - controlling Reporting Server browser interface access on OPSYS [80](#)
 - password encryption [97](#)
 - setting privileges [87](#)
 - SSL encryption [100](#), [103](#)
 - security LDAP [56](#)
 - security modes
 - LDAP security [56](#)
 - security DBMS [56](#)
 - security OPSYS [46](#)
 - security PTH [51](#)
 - security OPSYS [46](#)
 - security PTH [51](#)
 - selecting client code pages [598](#)
 - server activity [659](#), [663](#), [668](#), [674](#), [684](#), [768](#), [772](#)
 - server administration [17](#), [18](#), [573](#)
 - server agents [660](#), [668](#), [684](#), [690](#)
 - server configuration [17](#), [18](#), [26](#), [573](#)
 - server connections [674](#)
 - server deployment modes [580](#)
 - server deployment parameters [581](#)
 - server profiles [581](#)
 - setting [581](#)

- server diagnostics [757](#), [792](#)
 - EDAPRINT files [759](#)
 - record and playback [781](#), [787](#), [791](#)
 - trace files [768](#)
- server fine-tuning [704](#)
 - JVM memory usage [705](#)
- server management [17](#), [18](#), [573](#)
- server migration [658](#)
- server profiles
 - commands [32](#)
- server security [35](#)
- server trace files [772](#)
- service profiles [28](#)
- service properties [663](#)
 - agents [663](#)
- Session Log [773](#)
- SET parameters
 - CURRENCY_PRINT_ISO [596](#)
 - DATE_ORDER [596](#)
 - DATE_SEPARATOR [597](#)
 - TIME_SEPARATOR [597](#)
- SET TRACEUSER commands [769](#)
- setting edahlip variable for HLIPRINT output [767](#)
- setting properties [735](#)
 - deferred management, and [735](#)
- setting Reporting Server browser interface preferences [23](#)
- setting server security [35](#)
- setting user privileges [87](#)
- settings
 - application [620](#)
 - core engine [605](#)
 - email SMTP server [617](#)
 - locale [591](#)
 - migrate [658](#)
 - miscellaneous [622](#)
 - profile [618](#)
- Show General Privileges option [23](#)
- showing Session Log [773](#)
- sink machine [219](#)
- sort order with Unicode [830](#)
- special service statistics [691](#)
- special services statistics [690](#)
- SSL [100](#), [103](#)
- stack size [705](#)
- stacked commands [498](#)
- starting agents [680](#), [688](#)
- statistical variables [501](#)
- statistics [659](#)
 - agent [663](#)
 - agents [660](#), [668](#), [684](#), [690](#)
 - connection [675](#)
 - deferred requests [696](#)
 - for individual agents [668](#), [690](#)
 - for individual Java Service Agents [684](#)
 - Performance Report for Java Service Agents [682](#)
 - special services (FDS, SCHEDULER, DFM) [690](#), [691](#)

- statistics [659](#)
 - system-specific [668](#), [684](#), [690](#)
 - status statistics for Data Services [669](#)
 - stdlog log files [791](#)
 - stdout log files [791](#)
 - stopping agents [678](#)
 - by user id [679](#)
 - stopping connections [676](#)
 - stored procedures [455](#)
 - creating [456](#)
 - editing [459](#)
 - printing [459](#)
 - run-time processing [495](#)
 - running [471](#)
 - scheduling [488](#)
 - string manipulation and extraction functions [828](#)
 - string replacement functions [828](#)
 - SUBMIT subcommand [548](#)
 - supplying variable values [509](#), [511](#), [512](#)
 - Sybase ASE Adapter [823](#)
 - Unicode support [823](#)
 - synonym editor [275](#)
 - business view frame [285](#)
 - common options [279](#)
 - compute [299](#)
 - decode [292](#)
 - define [299](#)
 - expression calculator [301](#)
 - filters/groups frame [287](#)
 - joins [295](#)
 - synonym editor [275](#)
 - ribbon [275](#)
 - table/column pane [280](#)
 - synonym locations [182](#), [183](#)
 - synonym
 - enhancing [273](#)
 - generating [273](#)
 - uses [273](#)
 - synonyms [222](#), [245](#)
 - long names on z/OS [222](#)
 - system variables [501](#), [504](#), [608](#)
 - system-extended statistics [668](#)
- ## T
- temporary files [75](#), [220](#)
 - testing for variable values [515](#), [516](#)
 - text editor [459](#)
 - TIME_SEPARATOR parameter [597](#)
 - trace files [768](#), [772](#)
 - limiting size [771](#)
 - traces [772](#)
 - filtering [772](#)
 - in script file [773](#)
 - troubleshooting servers [757](#), [792](#)
 - EDAPRINT files [759](#)
 - record and playback [781](#), [787](#), [791](#)
 - trace files [768](#)
 - TSO commands [549](#)
 - tuning server performance [704](#)
 - JVM memory usage [705](#)

U

unconditional branching [513](#)

Unicode [819](#)

- access non-FOCUS data sources [821](#)
- Asian characters and [819](#)
- configuring WebFOCUS for UTF-8 [821](#)
- data type support for Relational adapters [825](#)
- Db2 data types [825](#)
- European characters and [819](#)
- Fixed files configuration [827](#)
- fonts for PDFs [831](#)
- manipulating characters [828](#)
- Microsoft SQL data types [825](#)
- Microsoft SQL Server configuration [823](#)
- Oracle configuration [821](#)
- Oracle data types [825](#)
- PDF output [831](#)
- reformatting characters [828](#)
- SAP BW configuration [826](#)
- SAP R/3-ECC configuration [826](#)
- sort order and [830](#)
- Sybase ASE Adapter [823](#)
- Sybase ASE data types [825](#)

universal character encoding standard [819](#)

Universal Naming Convention (UNC) [182](#)

UNIX language and code page values [598](#)

upload [256](#)

USE command [214](#), [219](#)

user actions [791](#)

user connections [579](#)

UTF (Unicode Transformation Format) [819](#)

V

variable values [509](#)

setting [511](#), [512](#)

testing [515](#), [516](#)

variables [501–504](#)

global [609](#)

naming [501](#), [502](#)

system [608](#)

version information [757](#)

W

web console appearance [590](#)

WebFOCUS Reporting Server [573](#)

running and monitoring [573](#)

Windows [182](#)

accessing a network drive [182](#)

language and code page values [598](#)

Workspace configuration settings [589](#)

Core Engine [605](#)

e-mail settings [617](#)

foccache settings [620](#)

General [625](#)

general options [592](#), [617](#), [618](#), [620](#), [622](#)

miscellaneous settings [622](#)

NLS (National Language Support) settings [592](#)

profile settings [618](#)

web console appearance [590](#)

Workspace Manager [18](#), [573](#)
 configuration privileges [574](#)
 configuring [573](#)
 monitoring connections [675](#)
 monitoring server activity [659](#), [663](#), [668](#),
 [674](#), [684](#), [689](#), [690](#), [772](#)

Z

z/OS commands [528](#)

