



IBM Net.Data for OS/400

Net.Data Administration and Programming Guide

Version 5





IBM Net.Data for OS/400

Net.Data Administration and Programming Guide

Version 5

Note

Before using this information and the product it supports, read the information in "Notices" on page 111.

October 2001 Edition

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

This edition applies to:

- IBM HTTP Server for AS/400 (Program 5769-DG1), Version 4 Release 4 Modification 0
- IBM HTTP Server for iSeries® (Program 5722-DG1), Version 5 Release 1 Modification 0

and to all subsequent versions and releases until otherwise indicated in new editions.

© Copyright IBM Corporation 1997, 2001.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Preface

Thank you for selecting Net.Data[®], the IBM[®] development tool for creating dynamic Web pages! With Net.Data you can rapidly develop Web pages with dynamic content by incorporating data from a variety of data sources and by using the power of programming languages you already know.

About Net.Data

With IBM's Net.Data product, you can create dynamic Web pages using data from both relational and non-relational database management systems (DBMSs), including DB2[®] and databases that can be accessed through DRDA[®], and using applications written in programming languages such as Java, JavaScript, C, C++, COBOL, and REXX.

Net.Data is a macro processor that executes as middleware on a Web server machine. You can write Net.Data application programs, called *macros*, that Net.Data interprets to create dynamic Web pages with customized content based on input from the user, the current state of your databases, other data sources, existing business logic, and other factors that you design into your macro.

A request, in the form of a URL (uniform resource locator), flows from a browser, such as Netscape Navigator or Internet Explorer, to a Web server that forwards the request to Net.Data for execution. Net.Data locates and executes the macro, and builds a Web page that it customizes based on functions that you write. These functions can:

- Encapsulate business logic within applications written in, but not limited to, C, C++, RPG, COBOL, Java, or REXX programming languages
- Access databases such as DB2
- Access other data sources such as flat files

Net.Data passes this Web page to the Web server, which in turn forwards the page over the network for display at the browser.

Net.Data can be used in server environments that are configured to use interfaces such as HyperText Transfer Protocol (HTTP) and Common Gateway Interface (CGI). HTTP is an industry-standard interface for interaction between a browser and Web server, and CGI is an industry-standard interface for Web server invocation of gateway applications like Net.Data. These interfaces allow you to select your favorite browser or Web server for use with Net.Data.

About this book

This book discusses administration and programming concepts for Net.Data, as well as how to configure Net.Data and its components, plan for security, and improve performance.

Building on your knowledge of programming languages and database, you learn how to use the Net.Data macro language to develop macros. You learn how to use Net.Data-provided language environments that access DB2 databases, as well as using RPG, and other programming languages to access your data.

This book may refer to products or features that are announced, but not yet available.

More information including sample Net.Data macros, demos, and the latest copy of this book, is available from the following World Wide Web site: <http://www.ibm.com/systems/i/software/netdata>.

Who should read this book

This book is intended for people involved in planning and writing Net.Data applications. To understand the concepts discussed in this book, you should be familiar with how a Web server works, understand simple SQL statements, and know HTML tags, including HTML form tags.

The Net.Data macro language, variables, and built-in functions, as well as operating system differences are described in *Net.Data Reference*.

About examples in this book

Examples used in this book are kept simple to illustrate specific concepts and do not show every way Net.Data constructs can be used. Some examples are fragments that require additional code to work.

Contents

Preface	iii
About Net.Data	iii
About this book	iii
Who should read this book	iv
About examples in this book.	iv
Chapter 1. Introduction	1
What is Net.Data?	1
Why use Net.Data?	2
Chapter 2. Configuring Net.Data	3
Copying the Net.Data executable to your CGI-BIN library	3
Customizing the Net.Data initialization file	4
Creating an initialization file	5
Configuration variable statements	5
Path configuration statements	15
Environment configuration statements	19
Setting up Net.Data language environments	20
Setting up the Java application language environment	20
Setting up the SQL language environment	21
Configuring the IBM HTTP Server.	21
Granting access rights to objects accessed by Net.Data	21
Chapter 3. Keeping your assets secure 23	
Using firewalls	23
Encrypting your data on the network.	25
Using authentication	25
Using authorization	26
Using Net.Data mechanisms.	26
Net.Data configuration variables	26
Macro Development Techniques	27
Chapter 4. Invoking Net.Data	31
Invoking Net.Data with a macro (macro request)	31
HTML links	32
HTML forms	33
Invoking a persistent macro	34
Persistent macro syntax	34
Examples	35
Chapter 5. Developing Net.Data macros 37	
Anatomy of a Net.Data macro	38
The DEFINE block	39
The FUNCTION block.	39
HTML blocks.	40
Net.Data macro variables.	42
Identifier Scope	42
Defining Variables	43
Referencing Variables	44
Variable types	45
Net.Data functions	51

Defining functions	52
Calling functions	55
Net.Data built-in functions	55
Generating document markup	59
HTML blocks.	59
Report blocks.	60
Conditional logic and looping in a macro	65
Conditional logic: IF blocks	65
Looping constructs: WHILE blocks	67
Chapter 6. Using language environments	69
Overview of Net.Data-supplied language environments.	70
Calling a language environment	70
Guidelines for handling error conditions.	70
Security	71
Direct Call language environment	72
Calling programs	72
Passing parameters to programs	72
Returning values from programs	74
Direct Call language environment example	74
Java Application language environment	76
Calling Java programs.	76
Passing parameters to Java programs.	76
Java Application language environment example	76
REXX language environment	78
Executing REXX programs	79
Passing parameters to REXX programs	79
REXX language environment example	80
SQL language environment	82
Executing SQL statements	82
Managing transactions in a Net.Data application	84
Managing multiple database connections	84
Stored procedures	85
SQL language environment example	89
System language environment	91
Issuing commands and calling programs	91
Passing parameters to programs	91
System language environment example	92

Chapter 7. Transaction management with persistent macros	93
About persistent macros	93
Defining a transaction	94
Starting a transaction	94
Specifying the macro HTML blocks in a transaction.	95
Ending a transaction	97
Defining the scope of a variable in a transaction	98
Specifying COMMIT and ROLLBACK in a transaction.	98
Example of a persistent macro	99

Chapter 8. Improving performance	103
---	------------

Net.Data caching of macros	103
Optimizing the language environments.	103
REXX language environment	103
SQL language environment.	104
System language environment.	104

Chapter 9. Serviceability features. 105

Net.Data trace log.	105
Net.Data error log.	105

Appendix A. Net.Data technical library 107

Appendix B. Net.Data sample macro 109

Notices	111
Trademarks	113

Glossary 115

Index 117

Chapter 1. Introduction

Net.Data is a server-side scripting language that extends Web servers by enabling the dynamic generation of Web pages using data from a variety of data sources. The data sources can include relational and non-relational database management systems such as DB2, DRDA-enabled databases, and flat file data. You can build applications rapidly using Net.Data's simple yet powerful scripting language. Net.Data allows reuse of existing business logic by supporting calls to applications written in a variety of programming languages, including Java, C/C++, RPG, CL, COBOL, REXX and others.

This chapter describes Net.Data and the reasons why you would choose to use it for your Web applications.

- "What is Net.Data?"
- "Why use Net.Data?" on page 2

What is Net.Data?

Using Net.Data macros, you can execute programming logic, access and manipulate variables, call functions, and use report-generating tools. A macro is a text file containing Net.Data language constructs, which are used to build an application that can consist of HTML, XML, Javascript, and language environment statements, such as SQL. Net.Data processes the macro to produce output that can be displayed by a Web browser. Macros combine the simplicity of HTML with the dynamic functionality of Web server programs, making it easy to add live data to static Web pages. The live data can be extracted from local or remote databases and from flat files, or be generated by applications and system services.

Figure 1 illustrates the relationship between Net.Data for OS/400®, the Web server, and supported data and programming language environments.

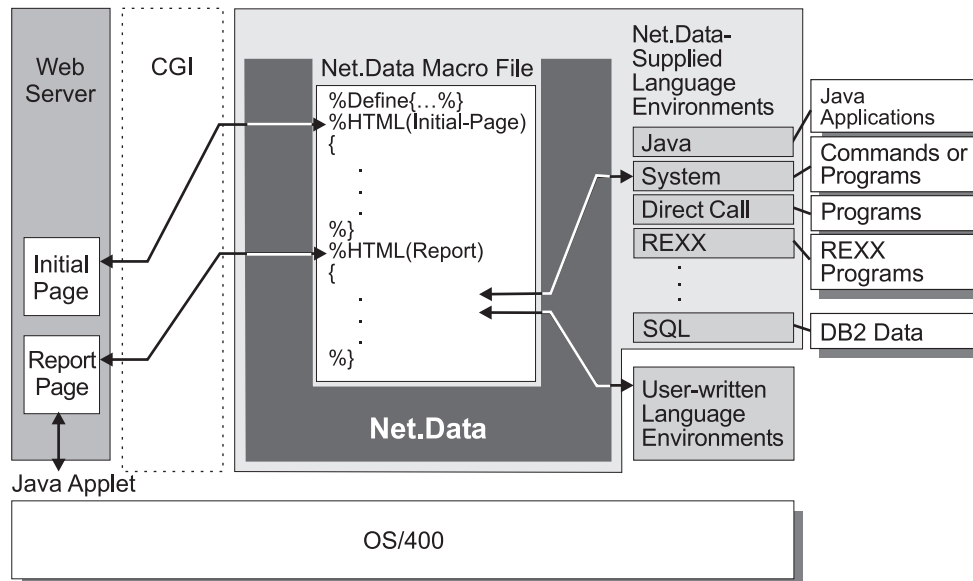


Figure 1. The Relationship between Net.Data for OS/400, the Web Server, and Supported Data and Program Sources

The Web server invokes Net.Data as a CGI application when it receives a URL that requests Net.Data services. The URL includes Net.Data-specific information, including the macro that is to be processed. When Net.Data finishes processing the request, it sends the resulting Web page to the Web server. The server passes it on to the Web client, where it is displayed by the browser.

Why use Net.Data?

Net.Data is a good choice for creating dynamic Web pages because using the macro language is simpler than writing your own Web server applications and because Net.Data lets you use languages that you already know, such as HTML, SQL, REXX, and JavaScript. Net.Data also provides language environments that access DB2 databases. In addition, changes to a macro can be seen instantaneously on a browser.

Net.Data complements data management capabilities that already exist on your operating system by enabling both data and related business logic for the Web. More specifically, Net.Data:

- Provides a simple, yet powerful macro language that allows for rapid development of Internet and Intranet applications.
- Permits the separation of data generation logic from presentation logic within your Web applications. Net.Data does not impose any restrictions on the method with which the data is presented (such as HTML or Javascript). This separation allows users to easily change the presentation of data using the latest presentation techniques.
- Allows you to use existing skills and business logic to generate Web pages by providing the ability to interface with programs written in C, C++, RPG, COBOL, REXX, Java or other languages.
- Provides the ability to develop complex Internet applications quickly, using a simple macro language.
- Provides high-performance access to data that is stored in DB2 and in any remote DRDA-enabled database.

Interpreted Macro Language

The Net.Data macro language is an interpreted language. When Net.Data is invoked to process a macro, Net.Data directly interprets each language statement in a sequential fashion, starting from the top of the file. Using this approach, any changes you make to a macro can be immediately seen when you next specify the URL that executes the macro. No recompilation is required.

Free Format

The Net.Data macro language has only a few rules about programming format. This simplicity provides programmers with freedom and flexibility. A single instruction can span many lines, or multiple instructions can be entered on a single line. Instructions can begin in any column. Spaces or entire lines can be skipped. Comments can be used anywhere.

Variables Without Type

Net.Data regards all data as character strings. Net.Data uses built-in functions to perform arithmetic operations on a string that represents a valid number, including those in exponential formats. Macro language variables are discussed in detail in “Net.Data macro variables” on page 42.

Built-in Functions

Net.Data supplies built-in functions that perform various processing, searching, and comparison operations for both text and numbers. Other built-in functions provide formatting capabilities and arithmetic calculations.

Error Handling

When Net.Data detects an error, messages with explanations are returned to the client. You can customize the error messages before they are returned to a user at a browser. See “Configuration variable statements” on page 5 and the *Net.Data Reference* for more information.

Chapter 2. Configuring Net.Data

Net.Data for OS/400 is delivered as a standard part of:

- IBM TCP/IP Connectivity Utilities/400 V3R2, V3R7, V4R1, and V4R2
- IBM HTTP Server for AS/400® V4R3 and subsequent releases

There is nothing extra to buy; and there is no Net.Data software that you need to download and install.

The AS/400 TCP/IP and HTTP Server software that you need comes standard with OS/400, but is optionally installed. The following optional software should be installed on your system for the following versions of the OS/400 operating system:

- For IBM OS/400 operating system Version 3 Release 2, Version 3 Release 7, and subsequent versions and releases (57xx-SS1):
 - IBM TCP/IP Connectivity Utilities/400 (57xx-TC1)
- For IBM OS/400 operating system Version 4 Release 3, and subsequent versions and releases (57xx-SS1):
 - IBM HTTP Server for AS/400 (57xx-DG1)

After installing Net.Data, complete the steps described in the following sections to configure Net.Data for OS/400. The steps include:

- “Copying the Net.Data executable to your CGI-BIN library”
- “Creating an initialization file” on page 5
- “Customizing the Net.Data initialization file” on page 4
- “Setting up Net.Data language environments” on page 20
- “Configuring the IBM HTTP Server” on page 21
- “Granting access rights to objects accessed by Net.Data” on page 21

Copying the Net.Data executable to your CGI-BIN library

Before using Net.Data, you must copy the Net.Data program object to the CGI-BIN library and provide access rights to the object.

To copy the Net.Data program object:

1. Using the Create Duplicate Object (CRTDUPOBJ) command, copy the Net.Data program object, DB2WWW, from the QHTTPSVR library to a CGI-BIN library.

Note: Prior to V4R3, the DB2WWW program object resided in library QTCP.

2. Change the DB2WWW program object in the CGI-BIN directory so that the user profile that CGI programs run under has access to the program object.

By default, the DB2WWW program object authority for *PUBLIC users is set to *EXCLUDE. To provide access to the program object, change the program object's authority for *PUBLIC users to *USE, or specifically give the user profile access to the DB2WWW program object.

You can copy the Net.Data program object to multiple libraries for different applications. This allows you to have more than one version of the Net.Data initialization file or multiple protection schemes. See “Customizing the Net.Data initialization file” on page 4 for more information about the Net.Data initialization file; see “Using authentication” on page 25 for information on authentication.

To copy the Net.Data program object to multiple libraries:

1. Copy the Net.Data program object, DB2WWW, to a library using the steps listed above.
2. Associate the Net.Data program object with a CL program in each library.
 - a. Create a CL program that calls the Net.Data program object located in the library specified in step 1.
 - b. Copy the CL program to each library.

In effect, the CL program you created becomes the Net.Data program object. If you do not associate the program object with a CL program, and copy the Net.Data program object DB2WWW to the different libraries, you get a -901 SQL code when using the SQL language environment.

In the following sections, the CL program you created should be treated as the Net.Data program object, if you chose to create the CL program to call Net.Data.

Customizing the Net.Data initialization file

Net.Data uses its initialization file to establish the settings of various configuration variables and to configure language environments and search paths. The settings of configuration variables control various aspects of Net.Data operation, such as the following:

- Specifying a SMTP server and a character set for sending e-mail
- Enabling the SQL language environment variable SHOWSQL

The language environment statements define the Net.Data language environments that are available and identify special input and output parameter values that flow to and from the language environments. The language environments enable Net.Data to access different data sources, such as DB2 databases and system services. The path statements specify the directory paths to files that Net.Data uses, such as macros and programs.

Creating the Net.Data initialization file is optional with Net.Data for OS/400. By using an initialization file, you can use shorter URLs and shorter references to programs and include files within your Net.Data macros. However, you are required to have an initialization file if you decide to create your own language environment.

If you do not create an initialization file, Net.Data runs as if you have configured an initialization file with only the supported language environment statements (see Chapter 6, “Using language environments,” on page 69 to learn about supported language environments). In this case, all macro, include, and executable references within the macro must be fully qualified.

The information contained in the initialization file is specified using three types of configuration statements, described in the following sections:

- “Configuration variable statements” on page 5
- “Path configuration statements” on page 15
- “Environment configuration statements” on page 19

See “Creating an initialization file” on page 5 to learn how to create an initialization file.

The sample initialization file shown below contains examples of these statements.

```
1 DTW SMTP_SERVER 9.5.5.78
2 MACRO_PATH /WWW/MACRO;/QSYS.LIB/WWW.LIB/MACRO.FILE
3 INCLUDE_PATH /WWW/MACRO;/QSYS.LIB/WWW.LIB/MACRO.FILE
4 EXEC_PATH /QSYS.LIB;/QSYS.LIB/WWW.LIB

5 ENVIRONMENT(MYLE1) /QSYS.LIB/LELIB.LIB/MYLE1.SRVPGM (IN VAR1, OUT VAR2)
```

Where:

- Line **1** sets the values of the configuration variables.

- Lines **2** to **4** define paths to the files that Net.Data needs to access.
- Line **5** specifies a user-defined ENVIRONMENT statement.

The text of each individual configuration statement must all be on one line. Ensure that the initialization file contains an ENVIRONMENT statement for each user-defined language environment that you call from your macros. If you fully qualify all references to files within the macro, you do not need to specify any of the path configuration statements.

The following sections describe how to create the initialization file and customize the configuration statements in the initialization file.

- “Configuration variable statements”
- “Path configuration statements” on page 15

Creating an initialization file

Creating an initialization file is optional when using Net.Data for OS/400. You should create an initialization file if:

- You want to set any of the Net.Data configuration variables to non-default values.
- You want to define the path statements for macro, include, and executable program files to shorten references to these files.
- You are using a language environment not supplied by Net.Data.

To create an initialization file:

1. In the library where the DB2WWW program object resides, use the Create Source Physical File (CRTSRCPF) command to create the initialization file.

File name:

INI

Member name:

DB2WWW

It is recommended that you create the initialization file with a record length of 240 because the text of configuration statements must all be on one line.

2. Use the Source Entry Utility (SEU) or a workstation editor to add configuration statements to the file as demonstrated in the sample macro and in the following sections.

If you create an initialization file and then update it, you do not need to end or restart the Web server in order for the changes to take effect. Net.Data reads the initialization file once, during the initial invocation by an HTTP server job. The configuration data is saved so that on subsequent Net.Data invocations, Net.Data does not have to read the initialization file. However, if a change is made to the initialization file, Net.Data detects the change to the initialization file and reads the initialization file again.

Authorization Tip: Ensure that the user IDs under which Net.Data executes have the appropriate access rights to this file. See “Granting access rights to objects accessed by Net.Data” on page 21 for more information.

Configuration variable statements

Net.Data configuration variable statements set the values of configuration variables. Configuration variables are used for various purposes. Some variables are required by a language environment to work properly or to operate in an alternate mode. Other variables control the character encoding or content of the Web page being constructed. Additionally, you can use configuration variable statements to define application-specific variables.

The configuration variables you use depend on the language environments you are using, as well as other factors that are specific to the application.

To update the configuration variable statements:

Customize the initialization file with the configuration variables that are required for your application. A configuration variable has the following syntax:

NAME [=] value-string

The equal sign is optional, as denoted by the brackets.

The following sub-sections describe the configuration variables statements that you can specify in the initialization file:

- "DTW_DEFAULT_ERROR_MESSAGE"
- "DTW_DEFAULT_MACRO" on page 7
- "DTW_ERROR_LOG_DIR" on page 7
- "DTW_ERROR_LOG_LEVEL" on page 7
- "DTW_JAVA_VMOPTIONS" on page 8
- "DTW_MACRO_CACHE_SIZE" on page 8
- "DTW_PAD_PGM_PARMS" on page 9
- "DTW_PROCESS_REPORT_ON_ERROR" on page 9
- "DTW_REMOVE_WS" on page 9
- "DTW_REPORT_BLOCK_VARIABLE_SCOPE" on page 10
- "DTW_RESTRICT_PATH_SEARCH" on page 10
- "DTW_SHOWSQL" on page 10
- "DTW_SMTP_CCSSID" on page 11
- "DTW_SMTP_CHARSET" on page 11
- "DTW_SMTP_SERVER" on page 12
- "DTW_SQL_ISOLATION" on page 13
- "DTW_SQL_NAMING_MODE" on page 14
- "DTW_TRACE_LOG_DIR" on page 14
- "DTW_TRACE_LOG_LEVEL" on page 14
- "DTW_TRACE_MERGE_RECORDS" on page 15
- "DTW_UPLOAD_DIR" on page 15

DTW_DEFAULT_ERROR_MESSAGE

Use the DTW_DEFAULT_ERROR_MESSAGE configuration variable to specify a generic error message for applications in production. This variable provides a generic message for error conditions that are not captured in any MESSAGE block.

If you still wish to see the actual error messages generated by Net.Data, use error message logging to capture the messages. See "Net.Data error log" on page 105 to learn about using the error log.

If the configuration variable is not specified, Net.Data displays its own provided message for the error condition.

Syntax:

DTW_DEFAULT_ERROR_MESSAGE [=] "message"

Example: Specifies a generic message

DTW_DEFAULT_ERROR_MESSAGE "This site is temporarily unavailable."

DTW_DEFAULT_MACRO

Use the DTW_DEFAULT_MACRO configuration variable to specify a default macro to be executed when Net.Data cannot find the macro requested in the URL.

You can use the environment variable PATH_INFO to determine the missing macro, as well as the HTML block that the URL requested.

If this configuration variable is not specified, Net.Data will print its own error message indicating that the macro could not be found.

Syntax:

DTW_DEFAULT_MACRO [=] name_of_macro

Example: Specifies the default macro that will execute.

DTW_DEFAULT_MACRO logurl.dtw

DTW_ERROR_LOG_DIR

Sets the directory where the error log is stored. See Chapter 9, "Serviceability features," on page 105 for more information on using this configuration variable.

Syntax:

DTW_ERROR_LOG_DIR [=] full_directory_path

DTW_ERROR_LOG_LEVEL

Sets the level of error logging. See Chapter 9, "Serviceability features," on page 105 for more information on using this configuration variable.

Syntax:

DTW_ERROR_LOG_LEVEL [=] OFF|INFORMATION|ERROR|INFORMATION+ERROR|ALL

Where:

OFF Specifies that no error messages are captured in the error message log. This is the default value.

INFORMATION

Net.Data will only log informational messages. For example, informational messages include messages that are issued when a URL is not valid, or when an HTML block is not found.

ERROR

Net.Data will only log error messages that are not handled by a message block.

INFORMATION+ERROR

Net.Data will log those messages that are logged by error log levels INFORMATION and ERROR.

ALL All messages are logged, including messages that are handled by message blocks.

Example:

DTW_ERROR_LOG_LEVEL ERROR

DTW_INTERPRET_HTTP_INPUT

Use the DTW_INTERPRET_HTTP_INPUT configuration variable to specify whether form or URL data coming in from a client should be interpreted.

When DTW_INTERPRET_HTTP_INPUT is set to YES, Net.Data will interpret data coming in from a client, which includes resolving variable references, which may reveal sensitive information if the data is

returned back to the client. Setting DTW_INTERPRET_HTTP_INPUT to NO will ensure that any variable references in form or URL data will not be resolved to a value.

Syntax:

```
DTW_INTERPRET_HTTP_INPUT [=] YES|NO
```

Where:

YES Specifies that form or URL data coming in from a client should be interpreted. This is the default value.

NO Specifies that form or URL data coming in from a client should not be interpreted.

Example:

```
DTW_INTERPRET_HTTP_INPUT NO
```

DTW_JAVA_VMOPTIONS

Specifies Java options that is to be applied to the Java environment created by the Java language environment.

Syntax:

```
DTW_JAVA_VMOPTIONS [=] java-options
```

Where:

java-options

Specifies java options to be used. The Java options that Net.Data recognizes are as follows:

- java.version
- java.home

When java.home is specified, Net.Data will set the JAVA_HOME environment variable in the job in which Net.Data is running in. java.home and java.version should not be specified at the same time. If both are specified, java.home will take precedence.

Example:

To run J2SE 5.0 32bit, specify java.home as follows:

```
DTW_JAVA_VMOPTIONS java.home=/QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit
```

To run Java classic 5.0, specify java.version as follows:

```
DTW_JAVA_VMOPTIONS java.version=1.5
```

DTW_MACRO_CACHE_SIZE

Indicates the memory size in megabytes that Net.Data should use when caching macros. When the cache size is exceeded, Net.Data removes old cached macros to make room in the cache. Net.Data removes the macros that have been used the least recently.

Syntax:

```
DTW_MACRO_CACHE_SIZE [=] size
```

Where:

size Specifies the cache memory size in number of megabytes. The default is 5 MB and caching is always enabled. If *size* is 0, no macros are cached. If *size* is 1 - 4, the default of 5 is used.

Example: Specifies a cache size of 16 MB.

```
DTW_MACRO_CACHE_SIZE 16
```

DTW_PAD_PGM_PARMS

Indicates to a language environment whether character parameters that are to be passed to a program or stored procedure are padded with blanks. Character parameters have a data type of CHARACTER or CHAR.

For IN or INOUT parameters, if the length of parameter value is less than the specified precision, blanks are inserted to the right of the parameter value until the length of the parameter value is the same as the precision.

For OUT parameters, the parameter value is set to *precision* blanks.

After the call to the program or stored procedure, all trailing blanks are removed from OUT and INOUT parameter values.

Set this variable in the Net.Data initialization file to specify a value for all of your macros. You can override the value by defining it in the macro. If DTW_PAD_PGM_PARMS is not defined in the macro, it uses the value in the initialization file.

DTW_PAD_PGM_PARMS is supported by the Direct Call and SQL language environments.

Syntax:

```
DTW_PAD_PGM_PARMS [=] YES|NO
```

Where:

- YES** Specifies that all IN and INOUT character parameter values are left justified and padded with blanks for the defined precision of the parameter before passing the parameters to a program or stored procedure. Trailing blanks are removed after the call to a program or stored procedure.
- NO** Specifies that no padding is added to character parameter values (values are NULL-terminated) when passing parameters to programs or stored procedures. Trailing blanks are not removed after calling a program or stored procedure. This is the default.

DTW_PROCESS_REPORT_ON_ERROR

Specifies whether the report block should be processed when an error occurs in a database function. This variable has been introduced in order to be compatible with Net.Data on other platforms, which do not process report blocks when an error occurs in a database function.

Syntax:

```
DTW_PROCESS_REPORT_ON_ERROR [=] YES|NO
```

Where:

- YES** Specifies that Net.Data will process the report block when an error occurs in a database function. This is the default.
- NO** Specifies that Net.Data will not process the report block when an error occurs in a database function.

DTW_REMOVE_WS

When this variable is set to YES, Net.Data removes extraneous white space from the HTML output. By compressing white space, this variable reduces the amount of data sent to the Web browser, thereby improving performance. The default is NO.

You can override this variable in the macro by using the DEFINE statement.

Syntax:

DTW_REMOVE_WS [=] YES|NO

DTW_REPORT_BLOCK_VARIABLE_SCOPE

The DTW_REPORT_BLOCK_VARIABLE_SCOPE configuration variable was introduced in order to fix the problem with report-block processing not honoring RPT_MAX_ROWS that was passed in as a parameter. Net.Data performed report block processing by getting the values for RPT_MAX_ROWS and START_ROW_NUM from the global scope (the SQL language environment honored local scope values, but report block processing still looked for RPT_MAX_ROWS in the global scope when processing the report block).

Note: Customers who do not have functions with parameter names of START_ROW_NUM and RPT_MAX_ROWS do not have to do anything. Those that do should consider using this support, but only after analyzing their application carefully.

Syntax:

DTW_REPORT_BLOCK_VARIABLE_SCOPE [=] GLOBAL|LOCAL

Where:

GLOBAL

Net.Data performs report block processing by getting the values for RPT_MAX_ROWS and START_ROW_NUM from the global scope (the SQL language environment will honor local scope values, but report block processing will still look for RPT_MAX_ROWS in the global scope when processing the report block). This value is the default.

LOCAL

When DTW_REPORT_BLOCK_VARIABLE_SCOPE is set to LOCAL and RPT_MAX_ROWS or START_ROW_NUM exists in the local scope (one can do this by defining parameters START_ROW_NUM or RPT_MAX_ROWS in the function definition), the values will be used. If the values of these variables is NULL or the NULL string, the default values for these variables will be used. The default value for START_ROW_NUM is 1 and the default value for RPT_MAX_ROWS is ALL. If the variables are not found in the local scope, the variable will be searched for in the global scope.

DTW_RESTRICT_PATH_SEARCH

Specifies that Net.Data should restrict file searches to directories specified in MACRO_PATH, EXEC_PATH and INCLUDE_PATH configuration variables.

Syntax:

DTW_RESTRICT_PATH_SEARCH [=] YES|NO

Where:

YES Specifies that Net.Data will restrict file searches to directories specified in MACRO_PATH, EXEC_PATH and INCLUDE_PATH configuration variables.

NO Specifies that if Net.Data does not find a file in directories specified in MACRO_PATH, EXEC_PATH and INCLUDE_PATH configuration variables, Net.Data will attempt to use file as-is to determine if it exists. This is the default.

DTW_SHOWSQL

Overrides the effect of setting SHOWSQL within your Net.Data macros.

Syntax:

DTW_SHOWSQL [=] YES|NO

Where:

- YES** Enables SHOWSQL in any macro that sets the value of SHOWSQL to YES.
- NO** Disables SHOWSQL in your macros, even if the variable SHOWSQL is set to YES. NO is the default.

Table 1 describes how the settings in the Net.Data initialization file and the macro determine whether the SHOWSQL variable is enabled or disabled for a particular macro.

Table 1. The Relationship Between Settings in the Net.Data Initialization File and the Macro for SHOWSQL

Setting of DTW_SHOWSQL	Setting SHOWSQL	SQL statement is displayed
NO	NO	NO
NO	YES	NO
YES	NO	NO
YES	YES	YES

DTW_SMTP_CC SID

Specifies the ASCII coded character set identifier (CCSID) associated with the Multi-purpose Internet Mail Extensions (MIME) character set specified in DTW_SMTP_CHARSET. The CCSID is to be used when translating data specified on the DTW_SENDMAIL function from EBCDIC to ASCII.

If DTW_SMTP_CC SID is specified, you must also specify DTW_SMTP_CHARSET. When specifying the CCSID, ensure that it is appropriate for the MIME character set specified in DTW_SMTP_CHARSET and that the CCSID is supported by the system. Table 2 lists common MIME character sets and the associated ASCII CCSID. If DTW_SMTP_CC SID is not set, Net.Data uses the CCSID associated with MIME character set ISO-8859-1, which is 819.

Syntax:

DTW_SMTP_CC SID [=] *ascii_ccsid*

where *ascii_ccsid* is the ASCII CCSID (a number between 1-65534) to be used when translating from EBCDIC to ASCII.

Example:

DTW_SMTP_CC SID 912

This ASCII CCSID corresponds to the MIME character set ISO-8859-2

DTW_SMTP_CHARSET

Specifies the Multi-purpose Internet Mail Extensions (MIME) character set that is to be used in the e-mail messages by the DTW_SENDMAIL function. If DTW_SMTP_CHARSET is specified, you must also specify DTW_SMTP_CC SID. When specifying the MIME character set, ensure that the character set is valid because Net.Data does not validate the value specified for this variable. If DTW_SMTP_CHARSET is not set, Net.Data uses the MIME character set ISO-8859-1, with the associated CCSID of 819.

Table 2 lists common MIME character sets and the associated ASCII CCSID.

Table 2. Character sets supported by Net.Data

MIME Standard Character Set	ASCII CCSID	Description
US-ASCII	367	US English
ISO-2022-JP	5052	Japan MBCS
ISO-8859-1	819	Latin-1

Table 2. Character sets supported by Net.Data (continued)

MIME Standard Character Set	ASCII CCSID	Description
ISO-8859-2	912	Latin-2
ISO-8859-5	915	Cyrillic
ISO-8859-6	1089	Arabic
ISO-8859-7	813	Greek
ISO-8859-8	916	Hebrew
ISO-8859-9	920	Latin-5

Syntax:

```
DTW_SMTP_CHARSET [=] character_set
```

Where *character_set* is the MIME character set to be used.

Example:

```
DTW_SMTP_CHARSET iso-8859-2
```

This MIME character set corresponds to the 912 ASCII CCSID.

DTW_SMTP_SERVER

Specifies the SMTP server to use for sending out e-mail messages using the DTW_SENDMAIL built-in function. The value of this variable can either be a host name or an IP address. If this variable is not set, Net.Data uses the local host as the SMTP server.

Syntax:

```
DTW_SMTP_SERVER [=] server_name
```

Where *server_name* is the host name or IP address of the the SMTP server that is to be used for sending e-mail messages.

Performance tip: Specify an IP address for this value to prevent Net.Data from connecting to a domain name server when retrieving the IP address of the specified SMTP server.

Example:

```
DTW_SMTP_SERVER 9.5.34.5
```

DTW_SQL_CONNECTION_IDLE_TIME

Indicates whether Net.Data should check for DB connections that have not been used for the specified amount of time. If DTW_SQL_CONNECTION_IDLE_TIME is specified, prior to performing the first SQL operation within a macro, Net.Data will determine if any active connections have not been used for the specified idle time. If there is, Net.Data will end all active connections.

Note: Net.Data will perform idle-time processing only if commitment control is active and transaction scope is multiple. For the case where commitment control is not active and/or transaction scope is single, connections can be recovered on-the-fly at the time of connection failure.

The use of DTW_SQL_CONNECTION_IDLE_TIME is primarily useful under the following conditions:

- Commitment control is active.
- The transaction scope is multiple.
- At least one remote database connection is established.

- The remote database management system where the remote database resides ends idle connections automatically after a period of idle time or the server where the database resides is rebooted periodically.

Syntax:

`DTW_SQL_CONNECTION_IDLE_TIME [=] maxIdleTime`

Where:

maxIdleTime

An integer value greater than zero representing the maximum amount of time, in minutes, an inactive connection is available before the connection and any active connections are closed. If *maxIdleTime* is less than 1, `DTW_SQL_CONNECTION_IDLE_TIME` is ignored.

Example: Specifies idle time of 8 minutes.

`DTW_SQL_CONNECTION_IDLE_TIME 8`

DTW_SQL_ISOLATION

The `DTW_SQL` language environment uses the `DTW_SQL_ISOLATION` configuration statement to determine the degree to which the database operations executed by the `DTW_SQL` language environment are isolated from concurrently executing processes.

Syntax:

`DTW_SQL_ISOLATION locking_method`

Where *locking_method* is one of the following values:

DTW_SQL_NO_COMMIT

Specifies not to use commitment control. For the OS/400 operating system, do not specify this value if a relational database is specified in the relational database directory and the relational database is on a non-OS/400 system.

DTW_SQL_READ_UNCOMMITTED

Specifies locking for the objects referred to in SQL ALTER, COMMENT ON, CREATE, DROP, GRANT, LABEL ON, and REVOKE statements and the rows updated, deleted, and inserted. The objects are locked until the end of the unit of work (transaction). Uncommitted changes in other processes can be seen.

DTW_SQL_READ_COMMITTED

Specifies locking for the objects referred to in SQL ALTER, COMMENT ON, CREATE, DROP, GRANT, LABEL ON, and REVOKE statements and the rows updated, deleted, and inserted. The objects are locked until the end of the unit of work (transaction). A row that is selected, but not updated, is locked until the next row is selected. Uncommitted changes in other processes cannot be seen.

DTW_SQL_REPEATABLE_READ

Specifies locking for the objects referred to in SQL ALTER, COMMENT ON, CREATE, DROP, GRANT, LABEL ON, and REVOKE statements and the rows selected, updated, deleted, and inserted. The objects are locked until the end of the unit of work (transaction). Uncommitted changes in other processes cannot be seen.

DTW_SQL_SERIALIZABLE

Specifies locking for the objects referred to in SQL ALTER, COMMENT ON, CREATE, DROP, GRANT, LABEL ON, and REVOKE statements and the rows selected, updated, deleted, and inserted. The objects are locked until the end of the unit of work (transaction). Uncommitted changes in other processes cannot be seen. All tables referred to in SELECT, UPDATE, DELETE, and INSERT statements are locked exclusively until the end of the unit of work (transaction).

DTW_SQL_NAMING_MODE

The DTW_SQL_NAMING_MODE configuration statement specifies how a table name can be specified in an SQL statement.

Syntax:

```
DTW_SQL_NAMING_MODE mode
```

Where *mode* is one of the following values:

SQL_NAMING

Specifies that tables are qualified by the collection name in the form:

```
collection.table
```

where *collection* is the name of the collection and *table* is the table name. The default qualifier is the user ID running the process that executes the SQL statement and is used when the table name is not explicitly qualified and the default collection name is not specified. SQL_NAMING is the default table name.

SYSTEM_NAMING

Specifies that files are qualified by library name in the form:

```
library/file
```

where *library* is the name of the library and *file* is the table name. The default search path is the library list (*LIBL) for the unqualified table name, if the table name (file) is not explicitly qualified and a default collection name (library) is not specified.

DTW_TRACE_LOG_DIR

Sets the directory where the trace log is stored. See Chapter 9, “Serviceability features,” on page 105 for more information on using this configuration variable.

Syntax:

```
DTW_TRACE_LOG_DIR [=] full_directory_path
```

Example:

```
DTW_TRACE_LOG_DIR /netdata/error/logs
```

DTW_TRACE_LOG_LEVEL

Sets the level of trace logging. See Chapter 9, “Serviceability features,” on page 105 for more information on using this configuration variable.

Syntax:

```
DTW_TRACE_LOG_LEVEL [=] OFF|APPLICATION|SERVICE
```

Where:

OFF Specifies that no trace data is captured in the trace log. This is the default value.

Application

Net.Data writes application-level trace messages to the trace log.

Service

Net.Data writes all trace messages to the trace log. This level of trace should only be used when asked to do so by IBM. The information in a SERVICE trace will not be helpful in debugging your own applications and will make reading the trace more difficult than necessary. You may be asked to set the level to SERVICE to help resolve product service questions.

Example:

DTW_TRACE_LOG_LEVEL APPLICATION

DTW_TRACE_MERGE_RECORDS

Specifies whether trace records from all threads should be merged into a single trace log file. See Chapter 9, “Serviceability features,” on page 105 for more information on using this configuration variable.

Syntax:

```
DTW_TRACE_MERGE_RECORDS [=] YES|NO
```

Where:

YES Specifies that trace records are written to one file, NETDATA.TRACE. This is the default value

NO Specifies that trace records are written to thread-specific file. The file name is NETDATA.TRACE.xxxxx, where xxxxx is the process/thread identifier.

Example:

```
DTW_TRACE_MERGE_RECORDS NO
```

DTW_UPLOAD_DIR

Specifies into which directory Net.Data will store files uploaded by the client. When this variable is not set, Net.Data will not accept the files for upload.

Syntax:

```
DTW_UPLOAD_DIR [=] path
```

Example:

```
DTW_UPLOAD_DIR /tmp/uploads
```

Path configuration statements

Net.Data determines the location of files and executable programs used by Net.Data macros from the settings of path configuration statements. The path statements are:

- “DTW_ATTACHMENT_PATH” on page 16
- “DTW_JAVA_CLASSPATH” on page 16
- “EXEC_PATH” on page 16
- “FFL_PATH” on page 17
- “INCLUDE_PATH” on page 17
- “MACRO_PATH” on page 18

These path statements identify one or more directories that Net.Data searches when attempting to locate macros, executable files, text files, and include files. The path statements that you need depend on the Net.Data capabilities that your macros use.

Update guidelines:

Some general guidelines apply to the path statements. Exceptions are noted in the description of each path statement.

- Separate each specified directory in the path statement with a semicolon (;).
- Forward slashes (/) and back slashes (\) are treated the same.
- Each path statement can specify multiple paths. Paths are searched from left to right in the order specified. This multiple-path capability lets you organize your files within multiple directories. For example, you can place each of your Web applications in its own directory.
- It is recommended to use absolute path statements.

The following sections describe the purpose and syntax of each path statement and provide examples of valid path statements.

DTW_ATTACHMENT_PATH

This path configuration statement specifies the path used to locate attachments to be sent using DTW_SENDMAIL.

Syntax:

```
DTW_ATTACHMENT_PATH [=] path
```

Example:

```
DTW_ATTACHMENT_PATH /usr/lpp/internet/server_root/pub/upload
```

DTW_JAVA_CLASSPATH

This path configuration statement specifies the path used to locate Java classes. Directories are separated by colons.

Syntax:

```
DTW_JAVA_CLASSPATH [=] path
```

Example: The following example shows the DTW_JAVA_CLASSPATH statement in the initialization file.

Net.Data initialization file:

```
DTW_JAVA_CLASSPATH /directory1/directory2:/QIBM/ProdData/Java400
```

EXEC_PATH

This path configuration statement identifies one or more directories that Net.Data searches for an external program that is invoked by the EXEC statement or an executable variable. If the program is found, the external program name is appended to the path specification, resulting in a fully qualified file name that is passed to the language environment for execution.

Syntax:

```
EXEC_PATH [=] path1;path2;...;pathn
```

Example: The following example shows the EXEC PATH statement in the initialization file and the EXEC statement in the macro that invokes an external program.

Net.Data initialization file:

```
EXEC_PATH /qsys.lib/programs.lib;/qsys.lib/rexx.lib/rexxpgms.file;
```

Net.Data macro:

```
%FUNCTION(DTW_REXX) myFunction() {  
  %EXEC{ myFunction.mbr %}  
%}
```

If the file myFunction.mbr is found in the /qsys.lib/rexx.lib/rexxpgms.file directory, the qualified name of the program is /qsys.lib/rexx.lib/rexxpgms.file/myFunction.mbr.

If the file is not found in the directories specified in the EXEC_PATH statement:

- If the specified path is absolute, Net.Data searches for the file in the specified path. For example, if the following EXEC statement were specified:

```
%EXEC{/qsys.lib/programs.lib/rpg1.pgm %}
```

Net.Data would search for the file rpg1.pgm in the /qsys.lib/programs.lib directory.

- If the specified path is relative, Net.Data searches the current working directory. For example, if the following EXEC statement were specified:

```
%EXEC { rpg1.pgm %}
```

then Net.Data would attempt to find the file rpg1.pgm in the current working directory.

FFI_PATH

This path configuration statement identifies one or more directories that Net.Data searches for a flat file that is referenced by a flat file interface (FFI) function.

Syntax:

```
FFI_PATH [=] path1;path2;...;pathn
```

Example: The following example shows an FFI_PATH statement in the initialization file.

Net.Data initialization file:

```
FFI_PATH /u/user1/ffi;/usr/lpp/netdata/ffi;
```

When the FFI language environment is called, Net.Data looks in the path specified in the FFI_PATH statement.

Because the FFI_PATH statement is used to provide security to those files not in directories in the path statement, there are special provisions for FFI files that are not found. See the FFI built-in functions section in *Net.Data Reference*.

INCLUDE_PATH

This path configuration statement identifies one or more directories that Net.Data searches, in the order in which they are specified, to find a file specified on an INCLUDE statement in a Net.Data macro. When it finds the file, Net.Data appends the include file name to the path specification to produce the qualified include file name.

Syntax:

```
INCLUDE_PATH [=] path1;path2;...;pathn
```

Example 1: The following example shows both the INCLUDE_PATH statement in the initialization file and the INCLUDE statement that specifies the include file.

Net.Data initialization file:

```
INCLUDE_PATH /u/user1/includes;/usr/lpp/netdata/includes
```

Net.Data macro:

```
%INCLUDE "myInclude.txt"
```

If the file *myInclude.txt* is found in the */u/user1/includes* directory, the fully-qualified name of the include file is */u/user1/includes/myInclude.txt*.

Example 2: The following example shows the INCLUDE_PATH statement and an INCLUDE file with a subdirectory name.

Net.Data initialization file:

```
INCLUDE_PATH /u/user1/includes;/usr/lpp/netdata/includes
```

Net.Data macro:

```
%INCLUDE "OE/oeheader.inc"
```

The include file is searched for in the directories `/u/user1/includes/0E` and `/usr/lpp/netdata/includes/0E`. If the file is found in `/usr/lpp/netdata/includes/0E`, the fully qualified name of the include file is `/usr/lpp/netdata/includes/0E/oeheader.inc`.

If the file is not found in the directories specified in the `INCLUDE_PATH` statement:

- If the specified path is absolute, Net.Data searches for the file in the specified path. For example, if the following `INCLUDE` statement were specified:

```
%INCLUDE "/u/user1/includes/oeheader.inc"
```

then Net.Data would search for the file `oeheader.inc` in the `/u/user1/includes` directory.

- If the specified path is relative, Net.Data searches the current working directory. For example, if the following `INCLUDE` statement were specified:

```
%INCLUDE "oeheader.inc"
```

then Net.Data would attempt to find the file `oeheader.inc` in the current working directory.

MACRO_PATH

This path configuration statement identifies the directories that Net.Data searches for Net.Data macros. For example, specifying the following URL requests the Net.Data macro with the path and file name `/macro/sqlm.dtw`:

```
http://server/cgi-bin/db2www/macro/sqlm.dtw/report
```

Syntax:

```
MACRO_PATH [=] path1;path2;...;pathn
```

The equal sign (=) is optional, as indicated by brackets.

Net.Data appends the path `/macro/sqlm.dtw/report` to the paths in the `MACRO_PATH` configuration statement, from left to right until Net.Data finds the macro. If the macro is not found, Net.Data will execute the macro defined for the `DTW_DEFAULT_MACRO` configuration variable, or it will print an error. See Chapter 4, "Invoking Net.Data," on page 31 for information on invoking Net.Data macros.

Example: The following example shows the `MACRO_PATH` statement in the initialization file and the related link that invokes Net.Data.

Net.Data initialization file:

```
MACRO_PATH /u/user1/macros;/usr/lpp/netdata/macros
```

HTML link:

```
<a href="http://server/cgi-bin/db2www/query.dtw/input">Submit another query.</a>
```

If the file `query.dtw` is found in the directory `/u/user1/macros`, then the fully-qualified path is `/u/user1/macros/query.dtw`.

If the file is not found in the directories specified in the `MACRO_PATH` statement, Net.Data searches for the file in the root (`/`) directory. For example, if the following URL is submitted:

```
http://myserver/cgi-bin/db2www/myfile.txt/report
```

and the file `myfile.txt` was not found in any of the directories specified in `MACRO_PATH`, then Net.Data attempts to find the file in the root (`/`) directory:

```
/myfile.txt
```

Environment configuration statements

An ENVIRONMENT statement configures a language environment. A language environment is a component of Net.Data that Net.Data uses to access a data source such as a DB2 database or to execute a program written in a language such as REXX. Net.Data provides a set of language environments, as well as an interface that allows you to create your own language environments. These language environments are described in Chapter 6, “Using language environments,” on page 69 and the language environment interface is described in *Net.Data Language Environment Interface Reference*.

Net.Data requires that an ENVIRONMENT statement for a particular language environment exist before you can invoke that language environment.

Net.Data for OS/400 does not require an ENVIRONMENT statement for language environments that are shipped with Net.Data. However, if a language environment statement is encountered, it overrides the default that Net.Data uses. It is recommended that ENVIRONMENT statements for Net.Data supplied language environments not be added to the Net.Data configuration file.

You can associate variables with a language environment by specifying the variables as parameters in the ENVIRONMENT statement. Net.Data implicitly passes the parameters that are specified on an ENVIRONMENT statement to the language environment as macro variables. To change the value of a parameter that is specified on an ENVIRONMENT statement in the macro, either assign a value to the variable using the DTW_ASSIGN() function or define the variable in a DEFINE section.

Important: If a variable is defined in a macro but is not specified on the ENVIRONMENT statement, the macro variable will not be passed to the language environment.

For example, a macro can define a DATABASE variable to specify the name of a database at which an SQL statement within a DTW_SQL function is to be executed. The value of DATABASE must be passed to the SQL language environment (DTW_SQL) so that the SQL language environment can connect to the designated database. To pass the variable to the language environment, you must add the DATABASE variable to the parameter list of the environment statement for DTW_SQL.

The sample Net.Data initialization file makes several assumptions about customizing the setting of Net.Data environment configuration statements. These assumptions may not be correct for your environment. Modify the statements appropriately for your environment.

To add or update an ENVIRONMENT statement:

ENVIRONMENT statements have the following syntax:

```
ENVIRONMENT(type) library_name (parameter_list, ...)
```

Parameters:

- *type*

The name by which Net.Data associates this language environment with a FUNCTION block that is defined in a Net.Data macro. You must specify the type of the language environment on a FUNCTION block definition to identify the language environment that Net.Data should use to execute the function.

- *library_name*

The name of the service program containing the language environment interfaces that Net.Data calls.

The service program name is specified with the .SRVPGM extension.

- *parameter_list*

The list of parameters that are passed to the language environment on each function call, in addition to the parameters that are specified in the FUNCTION block definition.

To set and pass the variables in the parameters list, define the variable in the macro.

You must define these parameters as configuration variables or as variables in your macro before executing a function that will be processed by the language environment. The following example specifies the variables in the ENVIRONMENT statement:

```
ENVIRONMENT(DTW_SQL) /QSYS.LIB/QTCP.LIB/QTMSQL.SRVPGM(IN
  DATABASE, LOGIN, PASSWORD, TRANSACTION_SCOPE)
```

If a function modifies any of its output parameters, the parameters keep their modified value after the function completes.

When Net.Data processes the initialization file, it does not load the language environment service programs. Net.Data loads a language environment service program when it first executes a function that identifies that language environment. The service program then remains loaded for as long as Net.Data is loaded.

Example: ENVIRONMENT statements for Net.Data-provided language environments

When customizing the ENVIRONMENT statements for your application, add the variables to the ENVIRONMENT statements that need to be passed from your initialization file to a language environment or that Net.Data macro writers need to set or override in their macros.

On OS/400, ENVIRONMENT statements are not required for Net.Data language environments and are not recommended. However, this example shows some of the default ENVIRONMENT statements that Net.Data uses.

```
1 MACRO_PATH      /WWW/MACRO;/QSYS.LIB/WWW.LIB/MACRO.FILE
2 INCLUDE_PATH   /WWW/MACRO;/QSYS.LIB/WWW.LIB/MACRO.FILE
3 EXEC_PATH      /QSYS.LIB;/QSYS.LIB/WWW.LIB

4 ENVIRONMENT(DTW_REXX) /QSYS.LIB/QHTTSPVR.LIB/QTMJREXX.SRVPGM ( )
5 ENVIRONMENT(DTW_SQL) /QSYS.LIB/QHTTSPVR.LIB/QTMSQL.SRVPGM (IN DATABASE,
  LOGIN, PASSWORD, SHOWSQL, TRANSACTION_SCOPE, DB_CASE,
  RPT_MAX_ROWS, START_ROW_NUM, DTW_SET_TOTAL_ROWS,
  OUT_DTWTABLE, SQL_CODE, TOTAL_ROWS)
```

Required: Each ENVIRONMENT statement must be on a single line.

Setting up Net.Data language environments

After you set the Net.Data configuration variables, some additional setup is required before the following language environments can function properly. The following sections describe the steps necessary to set up the language environments:

- “Setting up the Java application language environment”
- “Setting up the SQL language environment” on page 21

Setting up the Java application language environment

Before using the Java Application language environment, first introduced in OS/400 V4R4, complete the following steps:

1. Install the “AS/400 Developer Kit for Java” licensed program, product identifier 57xxJV1. The “AS/400 Developer Kit for Java” must be installed to run Java applications on the AS/400.
2. Set the DTW_JAVA_CLASSPATH path configuration variable in the Net.Data initialization file so Java can find the Java application classes. For more information on this path configuration statement, see “DTW_JAVA_CLASSPATH” on page 16.
3. Set the DTW_JAVA_VMOPTIONS configuration variable to the JDK level that is to be used by the Java Application Language Environment[®]. To learn more about this configuration variable, see “DTW_JAVA_VMOPTIONS” on page 8.

After setting up the Java Application language environment, see “Java Application language environment” on page 76 to learn how to use the Java Application language environment.

Setting up the SQL language environment

Before using the SQL language environment, complete the following steps:

1. Create a directory entry for the local database in the relational database directory (a directory entry with a remote location of *LOCAL), in addition to any remote databases that the SQL language environment needs to access.

Add the entry by using the Add Relational Database Directory Entry (ADDRDBDIRE) command.

If you are accessing a remote database, complete additional configuration steps, such as setting up communications between the local system and the remote system. For more information about distributed database support, see *IBM i Distributed Database Programming*.

2. Add or update configuration variables. The SQL language environment supports the following configuration variables that can be specified in a Net.Data initialization file:

DTW_SQL_ISOLATION

Determines the degree to which the database operations executed by the SQL language environment are isolated from concurrently executing processes

DTW_SQL_NAMING_MODE

Determines how a table name can be specified in an SQL statement

DTW_SHOWSQL

Enables the use of the macro variable SHOWSQL

To learn more about the Net.Data configuration variable statements, see “Configuration variable statements” on page 5.

After setting up the SQL language environment, see “SQL language environment” on page 82 to learn how to use the SQL language environment.

Configuring the IBM HTTP Server

The Common Gateway Interface (CGI) is an industry-standard interface that enables a Web server to invoke an application program such as Net.Data.

Configure the Web server to invoke Net.Data by adding the appropriate directives to the HTTP configuration file. For example, assuming the Net.Data program object resides in library CGI, then the following directives redirect Net.Data requests to /QSYS.LIB/CGI.LIB/DB2WWW.PGM:

```
.  
. .  
ScriptAlias /CGI-BIN/DB2WWW/ /QSYS.LIB/CGI.LIB/DB2WWW.PGM/  
<Directory /QSYS.LIB/CGI.LIB>  
  Options +ExecCGI +Includes  
  order allow,deny  
  allow from all  
</Directory>  
. . .
```

Granting access rights to objects accessed by Net.Data

Before using Net.Data, you need to ensure that the user IDs under which Net.Data executes have the appropriate access rights to objects that are referenced in a Net.Data macro and to the macro that a URL references.

More specifically, ensure that the user IDs under which Net.Data executes have the following authorizations:

- To read the Net.Data initialization file, INI.FILE/DB2WWW.MBR
- To execute the Net.Data executable files and service programs, and to search the directories (libraries) in the paths to the executable files and service programs
- To read the appropriate Net.Data macros and search the appropriate directories identified by the MACRO_PATH path configuration statement
- To execute the appropriate files and to search the appropriate directories identified by the EXEC_PATH path configuration statement
- To read the appropriate files and to search the appropriate directories identified by the INCLUDE_PATH path configuration statement
- To read and write the appropriate files, and to search the appropriate directories identified by the FFI_PATH path configuration statement
- To access any object that might be referenced by the target of a language environment statement. For example SQL language environment runs SQL statements, and SQL statements access database files, so the user ID that Net.Data is running under must have authority to the database files.

Examples:

Depending on the file system in which you choose to store your Net.Data macros, you need to authorize the user profile under which the Net.Data CGI program is run to the Net.Data macro. The following methods give the QTMHHTTP1 user profile authority (in V3R2 and V3R7, Internet Connection for AS/400 ran CGI programs only under the QTMHHTTP1 user profile.):

- In the root file system, use the Change Authority (CHGAUT) CL command to give authority to the user profile:

```
CHGAUT OBJ('/WWW') USER(QTMHHTTP1) DTAAUT(*RX)
CHGAUT OBJ('/WWW/macro') USER(QTMHHTTP1) DTAAUT(*RX)
CHGAUT OBJ('/WWW/macro/*') USER(QTMHHTTP1) DTAAUT(*RX)
```

You need to give authority to all objects in the path.

- In the library file system (QSYS.LIB), use the Grant Object Authority (GRTOBJAUT) CL command to give authority to the user profile:

```
GRTOBJAUT OBJ(WWW) OBJTYPE(*LIB) USER(QTMHHTTP1) AUT(*USE)
GRTOBJAUT OBJ(WWW/MACRO) OBJTYPE(*FILE) USER(QTMHHTTP1) AUT(*USE)
```

You need to give authority only to the library and the source physical file.

You can also use the CHGAUT CL command to give authority to objects in the QSYS.LIB file system as follows:

```
CHGAUT OBJ('/QSYS.LIB/WWW.LIB') USER(QTMHHTTP1) DTAAUT(*RX)
CHGAUT OBJ('/QSYS.LIB/WWW.LIB/MACRO.FILE') USER(QTMHHTTP1) DTAAUT(*RX)
```

Language environment-specific authority considerations are documented in each language environment section in Chapter 6, "Using language environments," on page 69.

Chapter 3. Keeping your assets secure

Internet security is provided through a combination of firewall technology, operating systems features, Web server features, Net.Data mechanisms, and the access control mechanisms that are part of your data sources.

You must decide on what level of security is appropriate for your assets. This chapter describes methods you can use for keeping your assets secure and also provides references to additional resources you can use to plan for the security of your Web site.

The following sections contain guidelines for protecting your assets. The security mechanisms described include:

- “Using firewalls”
- “Encrypting your data on the network” on page 25
- “Using authentication” on page 25
- “Using authorization” on page 26
- “Using Net.Data mechanisms” on page 26

For more information on protecting your assets, see the Internet security list of frequently asked questions (FAQ) at this Web site:

<http://www.w3.org/Security/Faq>

Using firewalls

Firewalls are collections of hardware, software, and policies that are designed to limit access to resources in a networked environment.

Firewalls:

- Protect the internal network from infiltration or intrusion
- Protect the internal network from data and programs that are brought in by internal users
- Limit internal user access to external data
- Limit the damage that can be done if the firewall is breached

Net.Data can be used with firewall products that execute in your environment.

The following possible configurations provide recommendations for managing the security of your Net.Data application. These configurations provide high-level information and assume that you have configured a firewall that isolates your secure intranet from the public Internet. Carefully consider these configurations with your organization's security policies:

- **High security configuration**

This configuration creates a subnetwork that isolates Net.Data and the Web server from both the secure intranet and the public Internet. The firewall software is used to create a firewall between the Web server and the public Internet, and another firewall between the Web server and the secured intranet, which contains DB2 Server. This configuration is shown by Figure 2 on page 24.

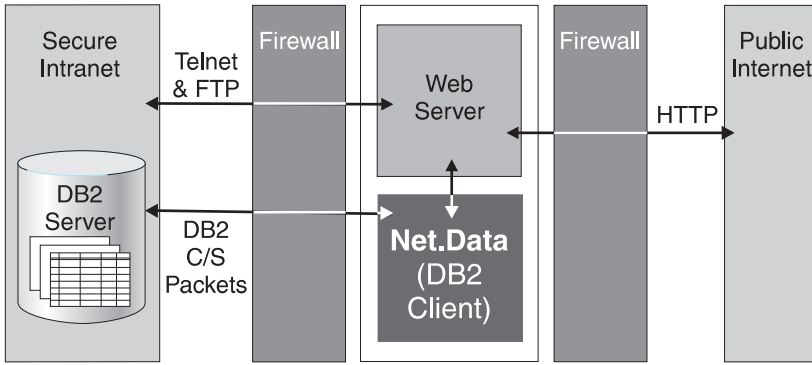


Figure 2. High Security Configuration

To set up this configuration:

- Install Net.Data on the Web server machine and ensure that Net.Data can access DB2 Server inside the intranet by configuring the firewall to allow DB2 traffic through the firewall. One method is to add a packet filtering rule to allow DB2 client requests from Net.Data and acknowledge packets from DB2 Server to Net.Data.
- Allow FTP and Telnet access between the Web server and the secure intranet. One method is to install a socks server on the Web server machine.
- In the packet filtering configuration file of the firewall software, specify that incoming TCP packets from the standard HTTP port can access the Web server. Also, specify that outgoing TCP acknowledge packets can go to any hosts on the public Internet from the Web server.

• **Intermediate security configuration**

In this configuration, firewall software isolates the secured intranet with DB2 server from the public Internet. Net.Data and the Web server are outside the firewall on a workstation platform. This configuration is simpler than the first, but still offers database protection. Figure 3 shows this configuration.

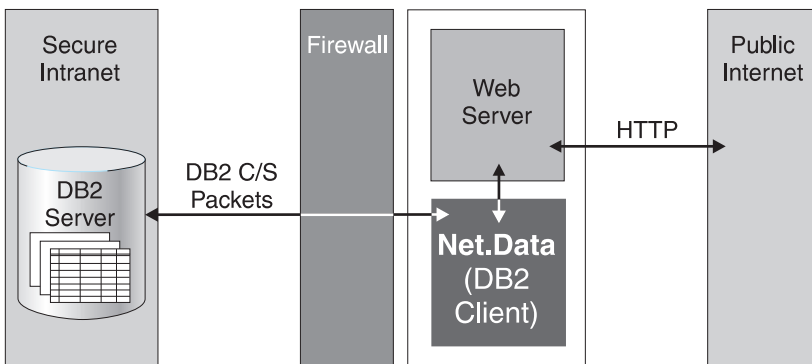


Figure 3. Intermediate Security Configuration:

The firewall must be configured to allow DB2 client requests to flow from Net.Data to DB2 and to allow acknowledge packets to flow from DB2 to Net.Data.

• **Low security configuration**

In this configuration, DB2 server and Net.Data are installed outside of the firewall and the secured intranet. They are not protected from external attacks. The firewall needs no packet filtering rules for this configuration. Figure 4 on page 25 shows this configuration.

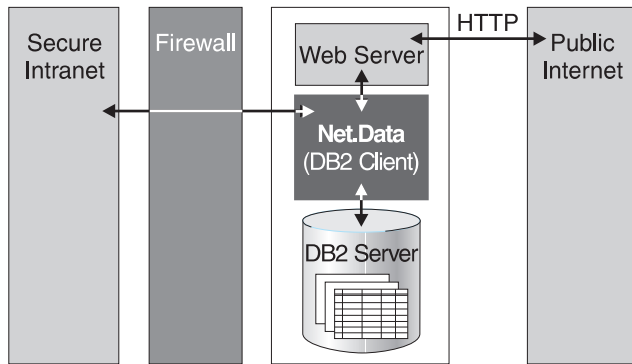


Figure 4. Low Security Configuration:

Encrypting your data on the network

You can encrypt all data that is sent between a client system and your Web server when you use a Web server that supports Secured Sockets Layer (SSL). This security measure supports the encryption of login IDs, passwords, and all data that is transmitted through HTML forms from the client system to the Web server and all data that is sent from the Web server to the client system.

Using authentication

Authentication is used to ensure that a user ID making a Net.Data request is authorized to access and update data within the application. Authentication is the process of matching the user ID with a password to validate that the request comes from a valid user ID. The Web server associates a user ID with each Net.Data request that it processes. The process or thread that is handling the request can then access any resource to which that user ID is authorized.

In an OS/400 environment, a user ID can become associated with the thread or process that is handling a Net.Data request in one of three ways:

Client-based authentication

The user is prompted for a local OS/400 user ID and password at the client. The Web server then authenticates the user. If successfully authenticated, the supplied user ID is associated with the request. Use of the special Web server `%%CLIENT%%` access control user ID enables this type of authentication.

Client-based authentication is supported by IBM's HTTP server starting in OS/400 V4R1.

Server-based authentication

The user ID of the Web server is associated with each request and the user is not prompted for a user ID or password. Use of the special Web server `%%SERVER%%` access control user ID enables this type of authentication.

By default, IBM's HTTP server runs CGI programs under the QTMHHTTP1 user ID (user profile). However, if the UserID directive is in effect or within a protection setup where the UserID subdirective has been specified, the program is executed under the specified user ID.

Surrogate authentication

A surrogate user ID that has the authority to access some predefined collection of resources is associated with the client request. This type of authentication requires the creation of surrogate user IDs with access authority that is appropriate for a group of users or class of requests. Authentication with surrogate user IDs usually uses validation list objects first introduced in V4R1. For more information and examples, see *OS/400 System API Reference*.

The approach that the Web server uses for associating a user ID with a client request is specified when the Web server is configured. For additional detail on access control, refer to your Web server documentation.

Tip: To protect Net.Data macros do the following:

1. Add protection directives in the Web server configuration file for the Net.Data program object.
2. Ensure the user ID that Net.Data will be running under has access rights to the macros. For more information on granting access rights, see “Granting access rights to objects accessed by Net.Data” on page 21.

Using authorization

Authorization provides a user with complete or restricted access to an object, resource, or function. Data sources such as DB2 provide their own authorization mechanisms to protect the information that they manage. These authorization mechanisms assume that the user ID associated with the Net.Data request has been properly authenticated, as explained in “Using authentication” on page 25. The existing access control mechanisms for these data sources then either permit or deny access based on the authorizations that are held by the authenticated user ID.

Using Net.Data mechanisms

In addition to the methods described above, you can use Net.Data configuration variables or macro development techniques to limit the activities of end users, to conceal corporate assets such as the design of your database, and to validate user-provided input values within production environments.

Net.Data configuration variables

Net.Data provides several configuration variables that can be used to limit the activities of end users or conceal the design of your database.

Control file access with path statements

Net.Data evaluates the settings of path configuration statements to determine the location of files and executable programs that are used by Net.Data macros. These path statements identify one or more directories that Net.Data searches when attempting to locate macros, executable files, include files, or other flat files. By selectively including directories on these path statements, you can explicitly control the files that are accessible by users at browsers. Refer to Chapter 2, “Configuring Net.Data,” on page 3 for additional detail about path statements.

You should also use authorization checking as described in “Using authorization” and verify that file names cannot be changed in INCLUDE statements as described in “Macro Development Techniques” on page 27.

Disable SHOWSQL for production systems

The SHOWSQL variable allows the user to specify that Net.Data displays the SQL statements specified within Net.Data functions at a Web browser. This variable is used primarily for developing and testing the SQL within an application and is not intended for use in production systems.

You can disable the display of SQL statements in production environments using one of the following methods:

- When using versions of Net.Data that support the DTW_SHOWSQL configuration variable, use this variable in the Net.Data initialization file to override the effect of setting SHOWSQL within your Net.Data macros. See “DTW_SHOWSQL” on page 10 for syntax and additional information.
- Use the DTW_ASSIGN() function as described in “Macro Development Techniques” on page 27.

See SHOWSQL in the variables chapter of *Net.Data Reference* for syntax and examples for the SHOWSQL Net.Data variable.

Set DTW_INTERPRET_HTTP_INPUT to NO

The DTW_INTERPRET_HTTP_INPUT configuration variable controls whether form or URL data coming in from a client should be interpreted. When DTW_INTERPRET_HTTP_INPUT is set to YES (the default), Net.Data will interpret data coming in from a client, which includes resolving variable references, which may reveal sensitive information if the data is returned back to the client. Setting DTW_INTERPRET_HTTP_INPUT to NO will ensure that any variable references in form or URL data will not be resolved to a value. See “DTW_INTERPRET_HTTP_INPUT” on page 7 for syntax and additional information.

Macro Development Techniques

Net.Data provides several mechanisms that allow users to assign values to input variables. To ensure that macros execute in the manner intended, these input variables should be validated by the macro. Your database and application should also be designed to limit a user's access to the data that the user is authorized to see.

Use the following development techniques when writing your Net.Data macros. These techniques will help you ensure that your applications execute as intended and that access to data is limited to properly authorized users.

Ensure that Net.Data variables cannot be overridden in a URL

The setting of Net.Data variables by a user within a URL overrides the effect of DEFINE statements used to initialize variables in a macro. This might alter the manner in which your macro executes. To safeguard against this possibility, initialize your Net.Data variables using the DTW_ASSIGN() function.

Example: Instead of using:

```
%define START_ROW_NUM = "1"
```

Use:

```
@DTW_ASSIGN(START_ROW_NUM, "1")
```

Assigning the variable this way prevents a query string assignment such as "START_ROW_NUM=10" from overriding your macro setting.

Validate that your SQL statements cannot be modified in ways that alter the intended behavior of your application

Adding a Net.Data variable to an SQL statement within a macro allows users to dynamically alter the SQL statement before executing it. It is the responsibility of the macro writer to validate user-provided input values and ensure that an SQL statement containing a variable reference is not being modified in an unexpected manner. Your Net.Data application should validate user-provided input values from the URL so the Net.Data application can reject invalid input. Your validation design process should include for the following steps:

1. Identify the syntax of valid input; for example, a customer ID must start with a letter and can contain only alphanumeric characters.
2. Determine what potential harm can be caused by allowing incorrect input, intentionally harmful input, or input entered to gain access to internal assets of the Net.Data application.
3. Include input verification statements in the macro that meet the needs of the application. Such verification depends on the syntax of the input and how it is used. In simpler cases it can be enough to check for invalid content in the input or to invoke Net.Data to verify the type of the input. If the syntax of the input is more complex, the macro developer might have to parse the input partially or completely to verify whether it is valid.

Example 1: Using the DTW_POS() string function to verify SQL statements

```
%FUNCTION(DTW_SQL) query1() {
    select * from shopper where shlogid = '${shlogid}'
%}
```

The value of the shlogid variable is intended to be a shopper ID. Its purpose is to limit the rows returned by the SELECT statement to rows that contain information about the shopper identified by the shopper ID. However, if the string "smith' or shlogid<>'smith" is passed as the value of the variable shlogid, the query becomes:

```
select * from shopper where shlogid = 'smith' or shlogid<>'smith'
```

This user-modified version of the original SQL SELECT statement returns the entire shopper table.

The Net.Data string functions can be used to verify that the SQL statement is not modified by the user in inappropriate ways. For example, the following logic can be used to ensure that single-quotes are not used to modify SQL statements:

```
@DTW_ADDQUOTE(shlogid, shlogid)
@query1()
```

The query then becomes:

```
select * from shopper where shlogid = 'smith'' or shlogid<>'smith'
```

Ensure that a file name in an INCLUDE statement is not modified in ways that alter the intended behavior of your application

If you specify the value for the file name with an INCLUDE statement using a Net.Data variable, then the file to be included is not determined until the INCLUDE file is executed. If your intent is to set the value of this variable within your macro, but to not allow a user at the browser to override the macro-provided value, then you should set the value of the variable using DTW_ASSIGN instead of DEFINE. If you do intend to permit the user at a browser to provide a value for the file name, then your macro should validate the value provided.

Example: A query string assignment such as filename=".././x" can result in the inclusion of a file from a directory not normally specified in the INCLUDE_PATH configuration statement. Suppose that your Net.Data initialization file contains the following path configuration statement:

```
INCLUDE_PATH /usr/lpp/netdata/include
```

and that your Net.Data macro contains the following INCLUDE statement:

```
%INCLUDE "${filename}"
```

A query string assignment of filename=".././x" would include the file /usr/lpp/x, which was not intended by the INCLUDE_PATH configuration statement specification.

The Net.Data string functions can be used to verify that the file name provided is appropriate for the application. For example, the following logic can be used to ensure that the input value associated with the file name variable does not contain the string "..":

```
@DTW_POS("..", ${filename}, result)
%IF (result > "0")
    %{ perform some sort of error processing %}
%ELSE
    %{ continue with normal processing %}
%ENDIF
```

Design your database and queries so that user requests do not have access to sensitive data about other users

Some database designs collect sensitive user data in a single table. Unless SQL SELECT requests are qualified in some fashion, this approach may make all of the sensitive data available to any user at a web browser.

Example: The following SQL statement returns order information for an order identified by the variable order_rn:

```

select setsstatcode, setsfailtype, mestname
  from merchant, setstatus
 where merfnbr = setsmenbr
 and   setsornbr = $(order_rn)

```

This method permits users at a browser to specify random order numbers and possibly obtain sensitive information about the orders of other customers. One way to safeguard against this type of exposure is to make the following changes:

- Add a column to the order information table that identifies the customer associated with the order information within a specific row.
- Modify the SQL SELECT statement to ensure that the SELECT is qualified by an authenticated customer ID provided by the user at the browser.

For example, if `shlogid` is the column containing the customer ID associated with the order, and `SESSION_ID` is a Net.Data variable that contains the authenticated ID of the user at the browser, then you can replace the previous SELECT statement with the following statement:

```

select setsstatcode, setsfailtype, mestname
  from merchant, setstatus
 where merfnbr = setsmenbr
 and   setsornbr = $(order_rn)
 and   shlogid = $(SESSION_ID)

```

Use Net.Data hidden variables

You can use Net.Data hidden variables to conceal various characteristics of your Net.Data macro from users that view your HTML source with their Web browser. For example, you can hide the internal structure of your database. See “Hidden variables” on page 47 for more information about hidden variables.

Request validation information from a user

You can create your own protection scheme based on user-provided input. For example, you can request validation information from a user through an HTML form and validate it using data that your Net.Data macro retrieves from a database or by calling an external program from a function defined in your Net.Data macro.

Chapter 4. Invoking Net.Data

This chapter describes how you invoke Net.Data using the Common Gateway Interface (CGI) Web server interface. See Chapter 2, “Configuring Net.Data,” on page 3 to learn more about configuring Net.Data for these interfaces.

Net.Data for OS/400 is invoked using Common Gateway Interface (CGI) and using a macro. This type of invocation method is called macro request. Additionally, you can invoke persistent macros, or macros that contain functions specifying transaction boundaries. For more information about persistent macros, see Chapter 7, “Transaction management with persistent macros,” on page 93

This chapter describes invoking Net.Data with a macro.

- “Invoking Net.Data with a macro (macro request)”
- “Invoking a persistent macro” on page 34

Invoking Net.Data with a macro (macro request)

A client browser invokes Net.Data by sending a request in the form of a URL. This section shows you how to invoke Net.Data by specifying a macro in the URL request.

The request sent to Net.Data has the following form.

```
http://server/Net.Data_invocation_path/filename/block[?name=val&...]
```

Parameters:

server Specifies the name and path of the Web server. If the server is the local server, you can omit the server name and use a relative URL.

Net.Data_invocation_path

The path and filename of the Net.Data executable file. For example, /cgi-bin/db2www/.

filename

Specifies the name of the Net.Data macro file. Net.Data searches for and tries to match this file name with the path statements defined in the MACRO_PATH initialization path variable. See “MACRO_PATH” on page 18 for more information.

block Specifies the name of the HTML block in the referenced Net.Data macro.

?name=val&...

Specifies one or more optional parameters passed to Net.Data.

You specify this URL directly in your browser. You can also specify it in an HTML link or build it using a form as follows:

- HTML link:
`any text`
- HTML form:
`<form method="method" ACTION="URL">any text</form>`

Parameters:

method Specifies the HTML method used with the form.

URL Specifies the URL used to run the Net.Data macro, the parameters of which are described above.

Examples

The following examples demonstrate the different methods of invoking Net.Data.

Example 1: Invoking Net.Data using an HTML link:

```
<a href="http://server/cgi-bin/db2www/myMacro.dtw/report">  
.  
.  
.  
</a>
```

Example 2: Invoking Net.Data using a form

```
<form method="post"  
  action="http://server/cgi-bin/db2www/myMacro.dtw/report">  
.  
.  
.  
</form>
```

Example 3: Invoking Net.Data macros in the qsys.lib file system, using an HTML link:

```
<a href="http://server/cgi-bin/db2www/myMacro.mbr/report">  
.  
.  
.  
</a>
```

Example 4: Invoking Net.Data macros in the qsys.lib file system, using a form:

```
<form method=post  
  action="http://server/cgi-bin/db2www/  
  qsys.lib/mylib.lib/myfile.file/myMacro.mbr/report">  
.  
.  
.  
</form>
```

The following sections describe HTML links and forms and more about how to invoke Net.Data with them:

- "HTML links"
- "HTML forms" on page 33

HTML links

If you are authoring a Web page, you can create an HTML link that results in the execution of an HTML block. When a user at a browser clicks on a text or image that is defined as an HTML link, Net.Data executes the HTML block within the macro.

To create an HTML link, use the HTML <a> tag. Decide which text or graphic you want to use as your hyperlink to the Net.Data macro, then surround it by the <a> and tags. In the HREF attribute of the <a> tag, specify the macro and the HTML block.

The following example shows a link that results in the execution of an SQL query when a user selects the text "List all monitors" on a Web page.

```
<a href="http://server/netdata-cgi/db2www/listA.d2w/report?hardware=mon">  
List all monitors</a>
```

Clicking on the link calls a macro named listA.dtw, which has an HTML block named "report", as in the following example:

```
%DEFINE DATABASE="MNS97"  
  
%FUNCTION(DTW_SQL) myQuery(){  
SELECT MODNO, COST, DESCRIP FROM EQPTABLE
```

```

WHERE TYPE='${hardware}'
%REPORT{
<h3>Here is the list you requested</h3>
%ROW{
<hr />
$(N1): $(V1), $(N2): $(V2)
<p>$(N3): $(V3)</p>
%}
%}
%}

%HTML (Report){
@myQuery()
%}

```

The query returns a table that contains model number, cost, and description information for each monitor that is described within the EQPTABLE table. The value of hardware in the SQL statement is taken from the URL input. See *Net.Data Reference* for a detailed description of the variables that are used in the ROW block.

HTML forms

You can dynamically customize the execution of your Net.Data macros using HTML forms. Forms allow users to provide input values that can affect the execution of the macro and the contents of the Web page that Net.Data builds.

The following example builds on the monitor list example in “HTML links” on page 32 by letting users at a browser use a simple HTML form to select the type of product for which information will be displayed.

```

<h1>Hardware Query Form</h1>
<hr>
<form method="post" action="/cgi-bin/db2www/listA.dtw/report">
<p>What type of hardware do you want to see?</p>
<ul>
<li><input type="radio" name="hardware" value="mon" checked /> Monitors</li>
<li><input type="radio" name="hardware" value="pnt" /> Pointing devices</li>
<li><input type="radio" name="hardware" value="prt" /> Printers</li>
<li><input type="radio" name="hardware" value="scn" /> Scanners</li>
</ul>

<input type="submit" value="submit" />
</form>

```

After the user at the browser makes a selection and clicks on the Submit button, the Web server processes the ACTION parameter of the FORM tag, which invokes Net.Data. Net.Data then executes the macro listA.dtw, which has an HTML block named "report" as listed above.

```

%DEFINE DATABASE="MNS97"

%FUNCTION(DTW_SQL) myQuery(){
SELECT MODNO, COST, DESCRIP FROM EQPTABLE
WHERE TYPE='${hardware}'
%REPORT{
<h3>Here is the list you requested</h3>
%ROW{
<hr />
$(N1): $(V1), $(N2): $(V2)
<p>$(N3): $(V3)</p>
%}
%}
%}

%HTML (Report){
@myQuery()
%}

```

In the above example, the value of `hdware` in the SQL statement is taken from the HTML form input. See *Net.Data Reference* for a detailed description of the variables that are used in the ROW block.

An input type that is given special treatment by `Net.Data` is the FILE input type. With this input type, users can upload a file to the server, which can be further processed by `Net.Data` or any other application on the server.

`Net.Data` tags the files with the correct codepage. The uploaded files are stored in the directory specified in `DTW_UPLOAD_DIR` and are given a unique name, determined using the following rules:

Syntax:

MacroFileName + '.' + *FormVarName* + '.' + *UniqueIdentifier* + '.' + *FormFileName*

MacroFileName

The name of the macro handling the request (the one called in the form). Only the filename is used, not the complete path.

FormVarName

The name of the variable used to identify the file in the form.

UniqueIdentifier

A string used to ensure uniqueness.

Example:

First, set `DTW_UPLOAD_DIR` in the `Net.Data` initialization file:

```
DTW_UPLOAD_DIR /tmp/uploads
```

Then, construct a form that invokes a macro and uses at least one input tag of type *file*.

```
<form method="post" enctype="multipart/form-data"
  action="/netdatadev/form.dtw/report">
  Name: <input type="text" name="name" /><br />
  Zip code: <input type="text" name="zipno" /><br />
  Resume: <input type="file" name="resume" /><br />
  <input type="submit" />
</form>
```

If a user were to submit the form, specifying the file `myresume.txt`, the resulting file would be written on the server with a name similar to:

```
/tmp/uploads/form.dtw.resume.20010108112341275-6245-021.myresume.txt
```

Invoking a persistent macro

This section shows you how to invoke persistent macros. These macros contain functions used for transaction processing. Invoking these macros is similar to regular macro request, in which you specify a server, macro, and HTML block. For persistent macros, you also specify a transaction handle, which identifies the HTML block as part of a transaction.

For more information about persistent macros and transaction processing, see Chapter 7, "Transaction management with persistent macros," on page 93.

Persistent macro syntax

Use the following syntax to invoke a persistent macro:

- HTML link:

```
<a href="http://server/Net.Data_invocation_path/transaction_handle/filename/
block/[?name=val&...]">any text</a>
```

- HTML form:

```
<form method=method ACTION="http://server/Net.Data_invocation_path/
transaction_handle/filename/block/
[?name=val&...]">any text</form>
```

- URL:

```
http://server/Net.Data_invocation_path/transaction_handle/filename/block/
[?name=val&...]
```

Parameters:

server Specifies the name of the Web server. If the server is the local server, you can omit the server name and use a relative URL.

Net.Data_invocation_path

The path and filename of the Net.Data executable file. For example, /cgi-bin/db2www/.

transaction_handle

Specifies which URLs are part of a transaction initiated by a Net.Data macro. This identifier is obtained by calling the DTW_RTVHANDLE built-in function and must follow the *Net.Data_invocation_path*.

filename

Specifies the name of the Net.Data macro. Net.Data searches for and tries to match this file name with the path statements defined in the MACRO_PATH initialization path variable. See "MACRO_PATH" on page 18 for more information.

block

Specifies the name of the HTML block in the referenced Net.Data macro.

method

Specifies the HTML method used with the form.

?name=val&...

Specifies one or more optional parameters passed to Net.Data.

Examples

The following examples demonstrate how to invoke persistent macros.

Example 1: A URL in a macro:

```
http://www.mycompany.com/cgi-bin/db2www/${handle}/mymacro.mac/report1
```

Example 2: A typical HTML block with links to other macro invocations that run in the same transaction

```
@DTW_STATIC()
```

```
...
```

```
%define handle = ""
@DTW_RTVHANDLE(handle)
```

```
%html(report) {
@DTW_ACCEPT(handle)
```

```
...
```

```
<a href="/cgi-bin/db2www/${handle}/qsys.lib/mylib.lib/macros.file/
pcgil.mbr/report2">continue</a><br />
```

```
<a href="/cgi-bin/db2www/${handle}/qsys.lib/mylib.lib/macros.file/
pcgil.mbr/quit">quit</a><br />
```

```
%}
```

Chapter 5. Developing Net.Data macros

A Net.Data macro is a text file consisting of a series of Net.Data macro language constructs that:

- Specify the layout of Web pages
- Define variables and functions
- Call functions that are built-in to Net.Data or defined in the macro
- Format the processing output and return it to the Web browser for display

The Net.Data macro contains two organizational parts: the declaration part and the presentation part, as shown in Figure 5.

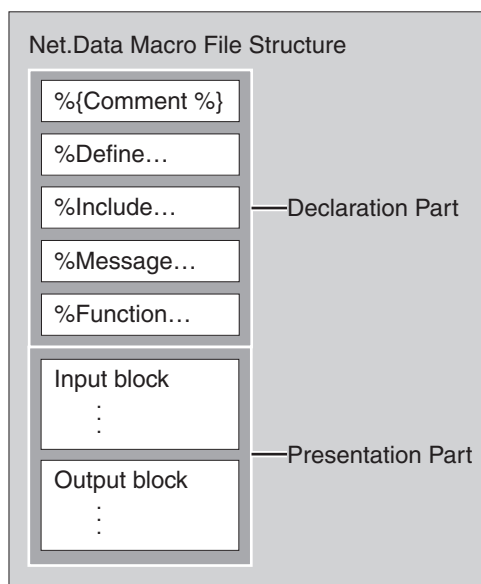


Figure 5. Macro Structure

- The *declaration part* contains the definitions of variables and functions in the macro.
- The *presentation part* contains HTML blocks that specify the layout of the generated Web document. The HTML blocks are made up of text presentation statements that are supported by your Web client, such as HTML, JavaScript, and well-formed XML.

You can use these parts multiple times and in any order. See *Net.Data Reference* for syntax of the macro parts and constructs.

Authorization Tip: Ensure that the user ID under which Net.Data executes has the authorization to read this file. See "Granting access rights to objects accessed by Net.Data" on page 21 for more information.

This chapter examines the different blocks that make up a Net.Data macro and methods you can use for writing the macro.

- "Anatomy of a Net.Data macro" on page 38
- "Net.Data macro variables" on page 42
- "Net.Data functions" on page 51
- "Generating document markup" on page 59
- "Conditional logic and looping in a macro" on page 65

Anatomy of a Net.Data macro

The macro consists of two parts:

- The declaration part, that contains definitions used in the presentation part. The declaration part uses two major optional blocks:
 - DEFINE block
 - FUNCTION block

The declaration part can also contain other language constructs and statements, such as EXEC statements, IF blocks, INCLUDE statements, and MESSAGE blocks. For more information about the language constructs, see the chapter about language constructs in *Net.Data Reference*.

Authorization Tip: Ensure that the user ID under which Net.Data executes has authorization to read and execute files referenced by EXEC statements and to read files referenced by INCLUDE statements. See “Granting access rights to objects accessed by Net.Data” on page 21 for more information.

- The presentation part defines the layout of the Web page, references variables, and calls functions using HTML blocks that are used as entry and exit points from the macro. When you invoke Net.Data, you specify a block name as an entry point for processing the macro. The HTML blocks are described in “HTML blocks” on page 40.

In this section, a simple Net.Data macro illustrates the elements of the macro language. This example macro presents a form that prompts for information to pass to a REXX program. The macro passes this information to an external REXX program called ompsamp\mbr, which echoes the data that the user enters. The results are then displayed on a second Web page.

First, look at the entire macro, and then each block in detail:

```
%{ ***** DEFINE block *****%}
%DEFINE {
  page_title="Net.Data Macro Template"
%}

%{ ***** FUNCTION Definition block *****%}
%FUNCTION(DTW_REXX) rexx1 (IN input) returns(result)
{
  %EXEC{ompsamp.mbr %}
%}

%FUNCTION(DTW_REXX) today () RETURNS(result)
{
  result = date()
%}

%{ ***** HTML Block: Input *****%}
%HTML (INPUT) {
<html>
<head>
<title>$(page_title)</title>
</head><body>
<h1>Input Form</h1>
Today is @today()

<form method="post" action="output">
Type some data to pass to a REXX program:
<input name="input_data" type="text" size="30" />
<p>
<input type="submit" value="enter" />
</p>
</form>

<hr>
<p>[<a href="/">Home page</a>]
</body></html>
```

```

%}

%{ *****          HTML Block: Output          *****%}
%HTML (OUTPUT) {
<html>
<head>
<title>$(page_title)</title>
</head><body>
<h1>Output Page</h1>
<p>@rexx1(input_data)
<p><hr>
<p>[<a href="/">Home page</a> |
<a href="input">Previous page</a>]
</body></html>
%}

```

The sample macro consists of four major blocks: the DEFINE, the FUNCTION, and the two HTML blocks. You can have multiple DEFINE, FUNCTION, and HTML blocks in one Net.Data macro.

The two HTML blocks contain text presentation statements such as HTML, which make writing Web macros easy. If you are familiar with HTML, building a macro simply involves adding macro statements to be processed dynamically at the server and SQL statements to send to the database.

Although the macro looks similar to an HTML document, the Web server accesses it through Net.Data using CGI. To invoke a macro, Net.Data requires two parameters: the name of the macro to process, and the HTML block in that macro to display.

When the macro is invoked, Net.Data processes it from the beginning. The following sections look at what happens as Net.Data processes the file.

The DEFINE block

The DEFINE block contains the DEFINE language construct and variable definitions used later in the HTML blocks. The following example shows a DEFINE block with one variable definition:

```

%{ *****          DEFINE Block          *****%}
%DEFINE {
    page_title="Net.Data Macro Template"
%}

```

The first line is a comment. A comment is any text inside %{ and %}. Comments can be anywhere in the macro. The next statement starts a DEFINE block. You can define multiple variables in one define block. In this example, only one variable, page_title, is defined. After it is defined, this variable can be referenced anywhere in the macro using the syntax, \$(page_title). Using variables makes it easy to make global changes to your macro later. The last line of this block, %}, identifies the end of the DEFINE block.

The FUNCTION block

The FUNCTION block contains declarations for functions invoked by the HTML blocks. Functions are processed by language environments and can execute programs, SQL queries, or stored procedures.

The following example shows two FUNCTION blocks. One defines a call to an external REXX program and the other contains inline REXX statements.

```

%{ *****          FUNCTION Block          *****%}
%FUNCTION(DTW_REXX) rexx1 (IN input) returns(result) { <-- This function accepts
                                                    one parameter and returns the
                                                    variable 'result', which is
                                                    assigned by the external REXX
                                                    program
%EXEC{ompsamp.mbr %} <-- The function executes an external REXX program
                        called "ompsamp.mbr"

```

```

%}
%FUNCTION(DTW_REXX) today () RETURNS(result) {
    result = date() <--- The single source statement for this function is
                        contained inline.
%}

```

The first function block, rexx1, is a REXX function declaration that in turn, runs an external REXX program called ompsamp.mbr. One input variable, input, is accepted by this function and automatically passed to the external REXX command. The REXX command also returns one variable called result. The contents of the result variable in the REXX command replaces the invoking @rexx1() function call contained in the OUTPUT block. The variables input and result are directly accessible by the REXX program, as shown in the source code for ompsamp.mbr:

```

/* REXX */
result = 'The REXX program received "'input'" from the macro.'

```

The code in this function echoes the data that was passed to it. You can format the resulting text any way you want by enclosing the requesting @rexx1() function call in normal mark-up style tags (like or). Rather than using the result variable, the REXX program could have written HTML tags to standard output using REXX SAY statements.

The second function block, also refers to a REXX program, today. However, the entire REXX program in this case is contained in the function declaration itself. An external program is not needed. Inline programs are allowed for REXX because it is an interpreted language that can be parsed and executed dynamically. Inline programs have the advantage of simplicity by not requiring a separate program file to manage. The first REXX function could also have been handled inline.

HTML blocks

HTML blocks define the layout of the Web page, reference variables, and call functions. HTML blocks are used as entry and exit points from the macro. An HTML block is always specified in the Net.Data macro request and every macro must have at least one HTML block.

The first HTML block in the example macro is named INPUT. The HTML(INPUT) contains the HTML for a simple form with one input field.

```

%{ ***** HTML Block: Input *****%}
%HTML (INPUT) { <--- Identifies the name of this HTML block.
<html>
<head>
<title>$(page_title)</title> <--- Note the variable substitution.
</head><body>
<h1>Input Form</h1>
Today is @today() <--- This line contains a call to a function.

<form method="post" action="output"> <--- When this form is submitted,
                                the "OUTPUT" HTML block is called.<p>
Type some data to pass to a REXX program:
<input name="input_data" <--- "input_data" is defined when the form
TYPE="text" SIZE="30" /> is submitted and can be referenced elsewhere in
                                this macro. It is initialized to whatever the
                                user types into the input field.

</p>
<input type="submit" value="enter" />

<hr>
<p>
[
<a href="/">Home page</a></p>
</body><html>
%} <--- Closes the HTML block.

```

The entire block is surrounded by the HTML block identifier, %HTML (INPUT) {...%}. INPUT identifies the name of this block. The name can contain underscores, periods, and any alphanumeric character; Net.Data does not distinguish by case. The HTML <title> tag contains an example of variable substitution. The value of the variable `page_title` is substituted into the title of the form.

This block also has a function call. The expression `@today()` is a call to the function `today`. This function is defined in the FUNCTION block that is described above. Net.Data inserts the result of the `today` function, the current date, into the HTML text in the same location that the `@today()` expression is located.

The ACTION parameter of the FORM statement provides an example of navigation between HTML blocks or between macros. Referencing the name of another block in an ACTION parameter accesses that block when the form is submitted. Any input data from an HTML form is passed to the block as implicit variables. This is true of the single input field defined on this form. When the form is submitted, data entered in this form is passed to the HTML(OUTPUT) block in the variable `input_data`.

You can access HTML blocks in other macros with a relative reference if the macros are on the same Web server. For example, the ACTION parameter `ACTION=" ../othermacro.dtw/main"` accesses the HTML block called `main` in the macro `othermacro.dtw`. Again, any data entered into the form is passed to this macro in the variable `input_data`.

When you invoke Net.Data, you pass the variable as part of the URL. For example:

```
<a href="/cgi-bin/db2www/othermacro.dtw/main?input_data=value">Next macro</a>
```

You can access or manipulate form data in the macro by referencing the variable name specified in the form.

The next HTML block in the example is the HTML(OUTPUT) block. It contains the HTML tagging and Net.Data macro statements that define the output processed from the HTML(INPUT) request.

```
%{ ***** HTML Block: Output *****%}
%HTML (OUTPUT) {
<html>
<head>
<title>$(page_title)</title> <--- More substitution.

</head><body>
<h1>Output Page</h1>
<p>@rexx1(input_data) <--- This line contains a call to function rexx1
passing the argument "input_data".

<p>
<hr>
<p>
[
<a href="/">Home page</a> |
<a href="input">Previous page</a>]
%}
```

Like the HTML(INPUT) block, this block is standard HTML with Net.Data macro statements to substitute variables and a function call. Again the `page_title` variable is substituted into the title statement. And, as before, this block contains a function call. In this case, it calls the function `rexx1` and passes to it the contents of the variable `input_data`, which it received from the form defined in the Input block. You can pass any number of variables to and from a function. The function definition specifies the number and the usage of variables that are passed.

Net.Data macro variables

Net.Data lets you define and reference variables in a Net.Data macro. In addition, you can pass these variables from the macro to the language environments and back. The variable names, values, and literal strings that are passed are called tokens. Net.Data puts no limit on the size of the tokens and will pass any token that the memory of your system can handle. Individual language environments, however, might provide restrictions on the token size.

Net.Data variables can be defined depending on the type of variable and whether it has a predefined value. These variables can be categorized into the following types, based on how they are defined:

- Explicitly defined variables using the DEFINE statement in the DEFINE block
- Predefined variables, which are variables that are made available by Net.Data and are set to a value. This value usually cannot be changed.
- Implicitly defined variables, which are of four types:
 - Variables that are not explicitly defined but are instantiated when first assigned a value.
 - Parameter variables that are part of a FUNCTION block definition and that can only be referenced within a FUNCTION block.
 - Variables that are instantiated by Net.Data and correspond to form data or query string data.
 - Variables that are associated with a Net.Data table and that can only be referenced within a ROW block or REPORT block.

The following sections describe:

- “Identifier Scope”
- “Defining Variables” on page 43
- “Referencing Variables” on page 44
- “Variable types” on page 45

Identifier Scope

If an identifier has global scope, then it can be referenced anywhere in a macro during a single request. The region where an identifier is visible is called its *scope*. The five types of scope are:

- Global
 - An identifier has global scope if you can reference it anywhere within a macro. Identifiers that have global scope are:
 - Net.Data built-in functions
 - Form data
 - Query string data
 - Variables instantiated from within an HTML block
- Macro
 - An identifier has this scope if its declaration appears outside of any block. A block starts with an opening bracket ({} and ends with a percent sign and bracket (%)). (DEFINE blocks are excluded from this definition.) Unlike an identifier with global scope, one with macro scope can only be referred to by items in the macro that follow the identifier's declaration.
- FUNCTION block or MACRO_FUNCTION block
 - An identifier has function block scope if:
 - The identifier is declared in the parameter list of the function definition.
 - If an identifier with the same name already exists outside the function definition, then Net.Data uses the identifier from the function parameter list within the function block.
 - The identifier is instantiated in the function block and is not declared or instantiated prior to the function call.

An identifier does not have function block scope if it has been declared or initialized outside of the function and is not declared in the function parameter list. The value of the identifier within the function block remains unchanged unless updated by the function.

- **REPORT block**

An identifier has report block scope if it can be referenced only from within a REPORT block (for example, table column names N_1, N_2, \dots, N_n). Only those variables that Net.Data implicitly defines as part of its table processing can have a report block scope. Any other variables that are instantiated have function block scope.

- **ROW block**

An identifier has row block scope if it can only be referenced from within a ROW block (for example, table value names V_1, V_2, \dots, V_n). Only those variables that Net.Data implicitly defines as part of its table processing can have a row block scope. Any other variables that are instantiated have function block scope.

Defining Variables

There are three ways to define variables in a Net.Data macro:

- Define statement or block
- HTML form tags
- Query string data

A variable value received from form or query string data overrides a variable value set by a DEFINE statement in a Net.Data macro.

- **DEFINE statement or block**

The simplest way to define a variable for use in a Net.Data macro is to use the DEFINE statement. The syntax is as follows:

```
%DEFINE variable_name="variable value"

%DEFINE variable_name={ variable value on multiple
                        lines of text %}

%DEFINE {
    variable_name1="variable value 1"
    variable_name2="variable value 2"
%}
```

The *variable_name* is the name you give the variable. Variable names must begin with a letter or underscore and can contain any alphanumeric character, an underscore, a period, or a hash (#). All variable names are case-sensitive, except *V_columnName*, which is a table variable.

For example:

```
%DEFINE reply="hello"
```

The variable *reply* has the value *hello*.

Two consecutive quotes alone is equal to an empty string. For example:

```
%DEFINE empty=""
```

The variable *empty* has an empty string.

If your variable contains special characters, such as an end-of-line, use block braces around the value:

```
%DEFINE introduction={
Hello,
My name is John.
%}
```

To include quotes in a string, you can use two quotes consecutively.

```
%DEFINE HI="say ""hello"""
```

You can also use block braces to escape the quotes:

```
%DEFINE HI={ say "hello" %}
```

To define several variables with one DEFINE statement, use a DEFINE block:

```
%DEFINE {  
    variable1="value1"  
    variable2="value2"  
    variable3="value3"  
    variable4="value4"  
%}
```

- **HTML form tags: SELECT, INPUT, and TEXTAREA**

You can use HTML FORM tags to assign values to variables, namely the SELECT, INPUT, and TEXTAREA tags. The following example uses standard HTML form tags to define Net.Data variables:

```
<input name="variable_name" TYPE=... />
```

or

```
<select name="variable_name">  
    <option>value one  
    <option>value two  
</select>
```

To assign a variable that spans multiple lines or contains special characters, such as quotes, the TEXTAREA tag can be used:

```
<textarea name="variable_name" ROWS="4">  
Please type the multi-line value  
of your variable here.  
</textarea>
```

The *variable_name* is the name you give the variable, and the value of the variable is determined from the input received in the form. See "HTML forms" on page 33 for an example of how this type of variable definition is used in a Net.Data macro.

- **Query string data**

You can pass variables to Net.Data through the query string. For example:

```
http://www.ibm.com/cgi-bin/db2www/stdqry1.dtw/input?field=custno
```

In the above example, the variable name, *field*, and the variable value, *custno*, specify additional data that Net.Data receives from the query string. Net.Data receives and processes the data as it would from form data.

Referencing Variables

You can reference a previously defined variable to return its value. To reference a variable in Net.Data macros, specify the variable name inside $\$($ and $)$. For example:

```
$(variableName)  
$(homeURL)
```

When Net.Data finds a variable reference, it substitutes the variable reference with the value of the variable. Variable references can contain strings, variable references, and function calls.

You can dynamically generate variable names. With this technique, you can use loops to process variably-sized tables or input data for lists that are built at run time, when the number in the list cannot be determined in advanced. For example, you can generate lists of HTML form elements that are generated based on records returned from an SQL query.

To use variables as part of your text presentation statements, reference them in the HTML blocks of your macro.

Invalid variable references: Invalid variable references are resolved to the empty string. For example, if a variable reference contains invalid characters such as an exclamation point (!), the reference is resolved to the empty string.

Valid variable names must begin with an alphanumeric character or an underscore, and they can consist of alphanumeric characters, including a period, underscore, and hash mark.

Example 1: Variable reference in a link

If you have defined the variable *homeURL*:

```
%DEFINE homeURL="http://www.ibm.com/"
```

You can refer to the home page as $\$(homeURL)$ and create a link:

```
<a href="\$(homeURL)">Home page</a>
```

You can reference variables in many parts of the Net.Data macro; check the language constructs in this chapter to determine in which parts of the macro variable references are allowed. If the variable has not yet been defined at the time it is referenced, Net.Data returns an empty string. A variable reference alone does not define the variable.

Example 2: Dynamically generated variable references

Assume that you run an SQL SELECT statement with any number of elements. You can create an HTML form with input fields using the following ROW blocks:

```
...
%ROW {
<input type=text name=@dtw_rconcat("I", ROW_NUM) size=10 maxlength=10 />
%}
...
```

Because you created INPUT fields, you would probably want to access the values that the user entered when the form is submitted to your macro for processing. You can code a loop to retrieve the values in a variable length list:

```
<pre>
...
@dtw_assign(rowIndex, "1")
%while (rowIndex <= rowCount) {
The value entered for row $(rowIndex) is: $(I$(rowIndex))
@dtw_add(rowIndex, "1", rowIndex) %}
...
</pre>
```

Net.Data first generates the variable name using the I(rowIndex)$ reference. For example, the first variable name would be I1. Net.Data then uses that value and resolves to the value of the variable.

Example 3: A variable reference with nested variable references and a function call

```
%define my = "my"
%define u = "lower"
%define myLOWERvar = "hey"
```

```
\$(my)@dtw_ruppercase(u)var)
```

The variable reference returns the value of hey.

Variable types

You can use the following types of variables in your macros.

- “Conditional variables” on page 46
- “Environment variables” on page 46
- “Executable variables” on page 46

- “Hidden variables” on page 47
- “List variables” on page 48
- “Table variables” on page 49
- “Miscellaneous variables” on page 49
- “Table processing variables” on page 50
- “Report variables” on page 50
- “Language environment variables” on page 51

If you assign strings to variables that are defined a certain way by Net.Data, such as ENVVAR, LIST, condition list variables, the variable no longer behaves in the defined way. In other words, the variable becomes a simple variable, containing a string.

See *Net.Data Reference* for syntax and examples of each type of variable.

Conditional variables

Conditional variables let you define a conditional value for a variable by using a method similar to an IF, THEN construct. When defining the conditional variable, you can specify two possible variable values. If the first variable you reference exists, the conditional variable gets the first value; otherwise the conditional variable gets the second value. The syntax for a conditional variable is:

```
varA = varB ? "value_1" : "value_2"
```

If varB is defined, varA="value_1", otherwise varA="value_2". This is equivalent to using an IF block, as in the following example:

```
%IF (varB)
  varA = "value_1"
%ELSE
  varA = "value_2"
%ENDIF
```

See “List variables” on page 48 for an example of using conditional variables with list variables.

Environment variables

You can reference environment variables that the Web server makes available to the process or thread that is processing your Net.Data request. When the ENVVAR variable is referenced, Net.Data returns the current value of the environment variable by the same name.

The syntax for defining environment variables is:

```
%DEFINE var=%ENVVAR
```

Where *var* is the name of the environment variable being defined.

For example, the variable SERVER_NAME can be defined as environment variable:

```
%DEFINE SERVER_NAME=%ENVVAR
```

And then referenced:

```
The server is $(SERVER_NAME)
```

The output looks like this:

```
The server is www.ibm.com
```

See *Net.Data Reference* for more information about the ENVVAR statement.

Executable variables

You can invoke other programs from a variable reference using executable variables.

Define executable variables in a Net.Data macro using the EXEC language construct in the DEFINE block. For more information about the EXEC language element, see the language constructs chapter in the *Net.Data Reference*. In the following example, the variable `runit` is defined to execute the executable program `testProg`:

```
%DEFINE runit=%EXEC "testProg"
```

`runit` becomes an executable variable.

Net.Data runs the executable program when it encounters a valid variable reference in a Net.Data macro. For example, the program `testProg` is executed when a valid variable reference is made to the variable `runit` in a Net.Data macro.

A simple method is to reference an executable variable from another variable definition. The following example demonstrates this method. The variable `date` is defined as an executable variable and `dateRpt` contains a reference to the executable variable.

```
%DEFINE date=%EXEC "date"
```

Wherever `$(date)` appears in the Net.Data macro, Net.Data searches for the executable program `date`, and when it locates it, displays:

```
Today is Tue 11-07-1999
```

When Net.Data encounters an executable variable in a macro, it looks for the referenced executable program using the following method:

1. It searches the directories specified by the `EXEC_PATH` in the Net.Data initialization file. See “`EXEC_PATH`” on page 16 for details.
2. If Net.Data does not locate the program, the system searches the directories defined by the system `PATH` environment variable or the library list. If it locates the executable program, Net.Data runs the program.

Restriction: Do not set an executable variable to the value of the output of the executable program it calls. In the previous example, the value of the variable `date` is `NULL`. If you use this variable in a `DTW_ASSIGN` function call to assign its value to another variable, the value of the new variable after the assignment is `NULL` also. The only purpose of an executable variable is to invoke the program it defines.

You can also pass parameters to the program to be executed by specifying them with the program name on the variable definition. In this example, the values of `distance` and `time` are passed to the program `calcMPH`.

```
%DEFINE mph=%EXEC "calcMPH $(distance) $(time)"
```

Hidden variables

You can use hidden variables to conceal the actual name of a variable from application users who view your Web page source with their Web browser. To define a hidden variable:

1. Define a variable for each string you want to hide, after the variable's last reference in the HTML block. Variables are always defined with the `DEFINE` language construct after they are used in the HTML block, as in the following example. The `$(variable)` variables are referenced and then defined.
2. In the HTML block where the variables are referenced, use double dollar signs instead of a single dollar sign to reference the variables. For example, `$(X)` instead of `$(X)`.

```
%HTML(INPUT) {  
<form ...>  
<p>Select fields to view:  
shanghai<select name="field">  
<option value="$(name)"> Name  
<option value="$(addr)"> Address  
...  
</form>  
%}
```

```
%DEFINE {
name="customer.name"
addr="customer.address"
%}
```

```
%FUNCTION(DTW_SQL) mySelect() {
SELECT $(Field) FROM customer
%}
```

...

When a Web browser displays the HTML form, \$\$ (name) and \$\$ (addr) are replaced with \$(name) and \$(addr) respectively, so the actual table and column names never appear on the HTML form.

Application users cannot tell that the true variable names are hidden. When the user submits the form, the HTML(REPORT) block is called. When @mySelect() calls the FUNCTION block, \$(Field) is substituted in the SQL statement with customer.name or customer.addr in the SQL query.

List variables

Use list variables to build a delimited string of values. They are particularly useful in helping you construct an SQL query with multiple items like those found in some WHERE or HAVING clauses. The syntax for a list variable is:

```
%LIST " value_separator " variable_name
```

Recommendation: The blanks are significant. Insert a space before and after the value separator for most cases. Most queries use Boolean or mathematical operators (for example, AND, OR, or >) for the value separator. The following example illustrates the use of conditional, hidden, and list variables:

```
%HTML(INPUT) {
<form method="post" action="/cgi-bin/db2www/example2.dtw/report">
<h2>Select one or more cities:</h2>
<input type="checkbox" name="conditions" value="$(cond1)" />Sao Paolo<br />
<input type="checkbox" name="conditions" value="$(cond2)" />Seattle<br />
<input type="checkbox" name="conditions" value="$(cond3)" />Shanghai<br />
<input type="submit" value="submit query" />
</form>
%}
```

```
%DEFINE{
DATABASE="custcity"
%LIST " OR " conditions
cond1="cond1='Sao Paolo'"
cond2="cond2='Seattle'"
cond3="cond3='Shanghai'"
whereClause= ? "WHERE $(conditions)"
%}
```

```
%FUNCTION(DTW_SQL) mySelect(){
SELECT name, city FROM citylist
$(whereClause)
%}
```

```
%HTML(REPORT){
@mySelect()
%}
```

In the HTML form, if no boxes are checked, conditions is empty, so whereClause is also empty in the query. Otherwise, whereClause has the selected values separated by OR. For example, if all three cities are selected, the SQL query is:

```
SELECT name, city FROM citylist
WHERE cond1='Sao Paolo' OR cond2='Seattle' OR cond3='Shanghai'
```

This example shows that Seattle is selected, which results in this SQL query:

```
SELECT name, city FROM citylist
WHERE cond1='Seattle'
```

Table variables

The table variable defines a collection of related data. It contains a set of rows and columns including a row of column headers. A table is defined in the Net.Data macro as in the following statement:

```
%DEFINE myTable=%TABLE(30)
```

The number following %TABLE is the limit on the number of rows that this table variable can contain. To specify a table with no limit on the number of rows, use the default or specify ALL, as shown in these examples:

```
%DEFINE myTable2=%TABLE
%DEFINE myTable3=%TABLE(ALL)
```

When you define a table, it has zero rows and zero columns. The only way you can populate a table with values is by passing it as an OUT or INOUT parameter to a function or by using the built-in table functions provided by Net.Data. The DTW_SQL language environment automatically puts the results of a SELECT statement into a table.

For non-database language environments, such as DTW_REXX or DTW_PERL, the language environment is also responsible for setting table values. However, the language environment script or program defines the table values cell-by-cell. See Chapter 6, “Using language environments,” on page 69 for more information about how language environments use table variables.

You can pass a table between functions by referring to the table variable name. The individual elements of the table can be referred to in a REPORT block of a function or by using the Net.Data table functions. See “Table processing variables” on page 50 for accessing individual elements in a table within a REPORT block, and see “Table functions” on page 58 for accessing individual elements of a table using a table function. Table variables are usually populated with values in an SQL function, and then used as input to a report, either in the SQL function or in another function after being passed to that function as a parameter. You can pass table variables as IN, OUT, or INOUT parameters to any non-SQL function. Tables can be passed to SQL functions only as OUT parameters.

If you reference a table variable, the contents of the table are displayed and formatted based on the setting of the DTW_HTML_TABLE variable. In the following example, the contents of myTable would be displayed:

```
%HTML (output) {
  $(myTable)
}
```

The column names and field values in a table are addressed as array elements with an origin of 1.

Miscellaneous variables

These variables are Net.Data-defined variables that you can use to:

- Affect Net.Data processing
- Find out the status of a function call
- Obtain information about the result set of a database query
- Determine information about file locations and dates

Miscellaneous variables can either have a predefined value that Net.Data determines or have values that you set. For example, Net.Data determines the DTW_CURRENT_FILENAME variable value based on the current file that it is processing, whereas you can specify whether Net.Data removes extra white space caused by tabulators and newline characters.

Predefined variables are used as variable references within the macro and provide information about the current status of files, dates, or the status of a function call. For example, to retrieve the name of the current file, you could use:

```
%REPORT {  
  <p>This file is <i>$(DTW_CURRENT_FILENAME)</i>.</p>  
}
```

Modifiable variable values are generally set using a DEFINE statement or with the @DTW_ASSIGN() function and let you affect how Net.Data processes the macro. For example, to specify whether white space is removed, you could use the following DEFINE statement:

```
%DEFINE DTW_REMOVE_WS="YES"
```

Table processing variables

Net.Data defines table processing variables for use in the REPORT and ROW blocks. Use these variables to reference values from SQL queries and function calls.

Table processing variables have a predefined value that Net.Data determines. These variables allow you to reference values from the result sets of SQL queries or function calls by the column, row, or field that is being processed. You can also access information about the number of rows being processed or a list of all the column names.

For example, as Net.Data processes a result set from an SQL query, it assigns the value of the variable Nn for each current column name, such that N1 is assigned to the first column, N2 is assigned to the second column, and so on. You can reference the current column name for your Web page output.

Use table processing variables as variable references within the macro. For example, to retrieve the name of the current column being processed, you could use:

```
%REPORT {  
  <p>Column 1 is <i>$(N1)</i>.</p>  
}
```

Table processing variables also provide information about the results of a query. You can reference the variable TOTAL_ROWS in the macro to show how many rows are returned from an SQL query, as in the following example:

```
Names found: $(TOTAL_ROWS)
```

Some of the table processing variables are affected by other variables or built-in functions. For example, TOTAL_ROWS requires that the DTW_SET_TOTAL_ROWS SQL language environment variable be activated so that Net.Data assigns the value of TOTAL_ROWS when processing the results from a SQL query or function call as in the following example:

```
%DEFINE DTW_SET_TOTAL_ROWS="YES"  
...
```

```
Names found: $(TOTAL_ROWS)
```

Report variables

Net.Data displays Web page output generated from the macro in a default report format. In an HTML block, the default report format displays a table using <pre> </pre> tags or using HTML table tags. You can override the default report by defining a REPORT block with instructions for displaying the output or by using one of the report variables to prevent the default report from being generated.

Report variables help you customize how your Web page output is displayed and how it is used with default reports and Net.Data tables. You must define these variables before using them with a DEFINE statement or with the @DTW_ASSIGN() function.

The report variables specify spacing, override default report formats, specify whether table output should be displayed in HTML or fixed-width characters, and specify other display features. For example, you can set DTW_HTML_TABLE to "yes" and Net.Data will generate the default report with HTML table tags instead of as a plain-text formatted table.

```
%DEFINE ALIGN="YES"  
...  
<p>Your query was on these columns: $(NLIST)
```

The START_ROW_NUM report variable lets you determine at which row to begin displaying the results of a query. For example, the following variable value specifies that Net.Data will begin displaying the results of a query at the third row.

```
%DEFINE START_ROW_NUM = "3"
```

You can also determine whether Net.Data uses HTML tags for default formatting. With DTW_HTML_TABLE set to YES, an HTML table is generated rather than a text-formatted table.

```
%DEFINE DTW_HTML_TABLE="YES"
```

```
%FUNCTION(DTW_SQL){  
SELECT NAME, ADDRESS FROM $(qTable)  
%}
```

Language environment variables

These variables are used with language environments and affect how the language environment processes a request.

With these variables, you can perform tasks such as establishing connections to databases, enabling NLS support, and determining whether the execution of an SQL statement is successful.

For example, you can use the SQL_STATE variable to access or display the SQL state value returned from the database.

```
%FUNCTION (DTW_SQL) val1() {  
  select * from customer  
%REPORT {  
  ...  
%ROW {  
  ...  
%}  
  SQLSTATE=$(SQL_STATE)  
%}
```

Net.Data functions

Net.Data provides built-in functions for use in your applications, such as word and string manipulation functions or functions that retrieve and set table variable functions. You can also define functions for use with your application, for example to call an external program or a stored procedure.

User-defined functions

Those functions that you define for use with your application, for example to call an external program or a stored procedure.

Net.Data built-in functions

Those functions that Net.Data provides for use in your applications, such as functions for manipulating words and strings and functions that get and set table variables.

These sections describe the following topics:

- "Defining functions" on page 52
- "Calling functions" on page 55
- "Net.Data built-in functions" on page 55

Defining functions

To define your own functions in the macro, use a FUNCTION block or MACRO_FUNCTION block:

FUNCTION block

Defines a subroutine that is invoked from a Net.Data macro and is processed by a language environment. FUNCTION blocks must contain language statements or calls to an external program.

MACRO_FUNCTION block

Defines a subroutine that is invoked from a Net.Data macro and is processed by Net.Data rather than a language environment. MACRO_FUNCTION blocks can contain any statement that is allowed in an HTML block.

Syntax: Use the following syntax to define functions:

FUNCTION block:

```
%FUNCTION(type) function-name([usage] [datatype] parameter, ...)
    [RETURNS(return-var)] {
    executable-statements
    [report-block]
    ...
    [report-block]
    [message-block]
%}
```

MACRO_FUNCTION block:

```
%MACRO_FUNCTION function-name([usage] parameter, ...) [RETURNS(return-var)] {
    executable-statements
    [report-block]
    ...
    [report-block]
%}
```

Where:

type Identifies a language environment that is configured in the initialization file. The language environment invokes a specific language processor (which processes the executable statements) and provides a standard interface between Net.Data and the language processor.

function-name

Specifies the name of the FUNCTION or MACRO_FUNCTION block. A function call specifies the *function-name*, preceded by an at (@) sign. See “Calling functions” on page 55 for details.

You can define multiple FUNCTION or MACRO_FUNCTION blocks with the same name so that they are processed at the same time. Each of the blocks must all have identical parameter lists. When Net.Data calls the function, all FUNCTION blocks with the same name or MACRO_FUNCTION blocks with the same name are executed in the order they are defined in the Net.Data macro.

usage Specifies whether a parameter is an input (IN) parameter, an output (OUT) parameter, or both types (INOUT). This designation indicates whether the parameter is passed into or received back from the FUNCTION block, MACRO_FUNCTION block, or both. The usage type applies to all of the subsequent parameters in the parameter list until changed by another usage type. The default type is IN.

datatype

The data type of the parameter. Some language environments expect data types for the parameters that are passed. For example, the SQL language environment expects them when calling stored procedures, as does the Direct Call language environment when calling programs. See Chapter 6, “Using language environments,” on page 69 to learn more about the supported data types for the language environment you are using.

parameter

The name of a variable with local scope that is replaced with the value of a corresponding argument specified on a function call. Parameters are passed to the language environment and are accessible to the executable statements using the natural syntax of that language or as environment variables. Parameter variable references are not valid outside the FUNCTION or MACRO_FUNCTION blocks.

return-var

Specify this parameter after the RETURNS keyword to identify a special OUT parameter. The value of the return variable is assigned in the function block, and its value is returned to the place in the macro from which the function was called. For example, in the following sentence, <p>My name is @my_name() ., @my_name() gets replaced by the value of the return variable. If you do not specify the RETURNS clause, the value of the function call is:

- NULL if the return code from the call to the language environment is zero
- The value of the return code, when the return code is non-zero.

executable-statements

The set of language statements that is passed to the specified language environment for processing after the variables are substituted and the functions are processed. *executable-statements* can contain Net.Data variable references and Net.Data function calls.

For FUNCTION blocks, Net.Data replaces all variable references with the variable values, executes all function calls, and replaces the function calls with their resulting values before the executable statements are passed to the language environment. Each language environment processes the statements differently. For more information about specifying executable statements or calling executable programs, see “Executable variables” on page 46.

For MACRO_FUNCTION blocks, the executable statements are a combination of text and Net.Data macro language constructs. In this case, no language environment is involved because Net.Data acts as the language processor and processes the executable statements.

report-block

Defines one or more REPORT blocks for handling the output of the FUNCTION or MACRO_FUNCTION block. See “Report blocks” on page 60.

message-block

Defines the MESSAGE block, which handles any messages for error conditions returned by the FUNCTION block. For more information on how to capture error conditions, see “Message blocks.”

Define functions outside of any other block and before they are called in the Net.Data macro.

Using special characters in functions

When characters that match Net.Data language constructs syntax are used in the language statements section of a function block as part of syntactically valid embedded program code (such as REXX or Perl), they can be misinterpreted as Net.Data language constructs, causing errors or unpredictable results in a macro.

Use one of the following methods to use COMMENT block delimiter characters, or any other Net.Data special characters as part of your embedded program code, without having them interpreted by Net.Data as special characters:

- Use the EXEC statement to call the program code, rather than putting the code inline.
- Use a variable reference to specify the special characters.

Message blocks

The MESSAGE block lets you determine how to proceed after a function call, based on the success or failure of the function call, and lets you display information to the caller of the function. When processing

a message, Net.Data sets the language environment variable RETURN_CODE for each function call to a FUNCTION block. RETURN_CODE is not set on a function call to a MACRO_FUNCTION block.

A MESSAGE block consists of a series of message statements, each of which specifies a return code value, message text, and an action to take. The syntax of a MESSAGE block is shown in the language constructs chapter of the *Net.Data Referencebook*.

A MESSAGE block can have a global or a local scope. If it is specified at the outermost macro layer, the MESSAGE block has global scope and is active for all function calls executed in the Net.Data macro. If you define more than one global MESSAGE block, the last one defined is active. However, if the MESSAGE block is defined in a FUNCTION block, its scope is local to that FUNCTION block (except for Net.Data built-in functions, whose errors are handled by global message blocks).

Net.Data uses these rules to process the value of the RETURN_CODE or SQL_STATE variables from a function call:

1. Check the local MESSAGE block for an exact match of the value of the RETURN_CODE or SQL_STATE; exit or continue as specified.
2. If the value is not 0, check local MESSAGE block for +default or -default; depending on the sign of the value, exit or continue as specified.
3. If the value is not 0, check local MESSAGE block for default; exit or continue as specified.
4. Check global MESSAGE block for an exact match of the RETURN_CODE or SQL_STATE; exit or continue as specified.
5. If the value is not 0, check global MESSAGE block for +default or -default; depending on the sign of the value, exit or continue as specified.
6. If the value is not 0, check global MESSAGE block for default; exit or continue as specified.
7. If the value is not 0, issue Net.Data internal default message and exit.

The following example shows part of a Net.Data macro with a global MESSAGE block and a MESSAGE block for a function.

```
%{ global message block %}
%MESSAGE {
  -100      : "Return code -100 message"   : exit
   100      : "Return code 100 message"    : continue
  +default : {
This is a long message that spans more
than one line. You can use HTML tags, including
links and forms, in this message. %} : continue
%}

%{ local message block inside a FUNCTION block %}
%FUNCTION(DTW_REXX) my_function() {
  %EXEC { my_command.mbr %}
  %MESSAGE {
    -100      : "Return code -100 message"   : exit
     100      : "Return code 100 message"    : continue
    -default : {
This is a long message that spans more
than one line. You can use HTML tags, including
links and forms, in this message. %} : exit
  %}
}
```

If *my_function()* returns with a RETURN_CODE value of 50, Net.Data processes the error in this order:

1. Check for an exact match in the local MESSAGE block.
2. Check for +default in the local MESSAGE block.
3. Check for default in the local MESSAGE block.
4. Check for an exact match in the global MESSAGE block.

5. Check for `+default` in the global MESSAGE block.

When Net.Data finds a match, it sends the message text to the Web browser and checks the requested action.

When you specify `continue`, Net.Data continues to process the Net.Data macro after printing the message text. For example, if a macro calls `my_functions()` five times and error 100 is found during processing with the MESSAGE block in the example, output from a program can look like this:

```
.  
. .  
11 May 1997          $245.45  
13 May 1997          $623.23  
19 May 1997          $ 83.02  
return code 100 message  
22 May 1997          $ 42.67  
  
Total:              $994.37
```

Calling functions

Use a Net.Data function call statement to call both user-defined functions and built-in functions. Use the at (@) character followed by a function name or a macro function name:

```
@function_name([ argument,... ])
```

function_name

This is the name of the function or macro function to invoke. The function must already be defined in the Net.Data macro, unless this is a built-in function.

argument

This is the name of a variable, a quoted string, a variable reference, or a function call. Arguments on a function call are matched up with the parameters on a function or macro function parameter list. And, each parameter is assigned the value of its corresponding argument while the function or macro function is being processed. The arguments must be the same number and type as the corresponding parameters.

Quoted strings as arguments can contain variable references and functions calls.

Example 1: Function call with a text string argument

```
@myFunction("abc")
```

Example 2: Function call with a variable and a function call arguments

```
@myFunction(myvar, @DTW_rADD("2","3"))
```

Example 3: Function call with a text string argument that contains a variable reference and a function call

```
@myFunction("abc$(myvar)def@DTW_rADD("2","3")ghi")
```

Net.Data built-in functions

Net.Data provides a large set of built-in functions to simplify Web page development. These functions are already defined by Net.Data, so you do not need to define them. You can call these functions as you would call other functions.

Figure 6 on page 56 shows how the Net.Data built-in functions and the macro interact.

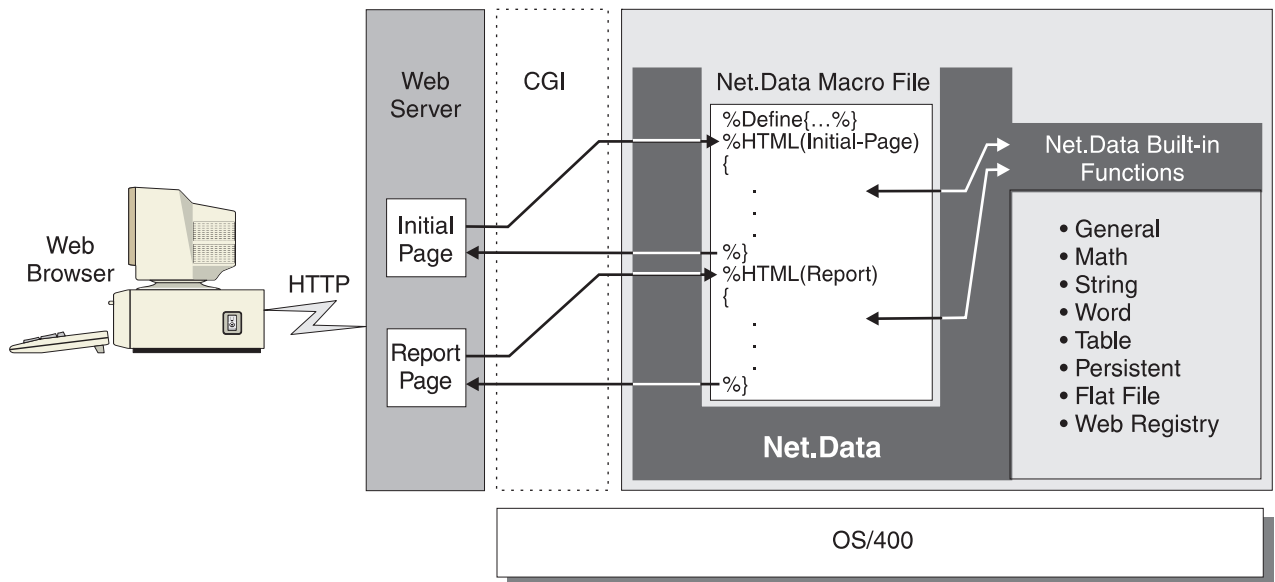


Figure 6. Net.Data Built-in Functions

Built-in functions can return their results in three ways, depending on its prefix:

- **DTW_**, **DTWF_**, and **DTWR_**: The results of the call are returned in an output parameter or no result is returned. (**DTWF_** is the prefix for flat file functions. **DTWR_** is the prefix for Web registry functions.)
- **DTW_r** and **DTWR_r**: The results of the function call replace the function call in the macro, in the same way the value of the RETURNS keyword replaces the function call for a user-defined function which has specified a RETURNS keyword.
- **DTW_m**: Multiple results are returned in each of the parameters passed to the function.

Some built-in functions do not have each type. To determine which type a particular built-in function has, see the Net.Data built-in functions chapter in *Net.Data Reference*.

The following sections provide a high-level overview of the Net.Data built-in functions. Use these functions to perform general purpose, math, string, word, or table manipulation functions. Some of these functions require variables to be set prior to their use or must be used in a specific context. See *Net.Data Reference* for descriptions of each function with syntax and examples.

- “General purpose functions”
- “Math functions” on page 57
- “String functions” on page 57
- “Word functions” on page 57
- “Table functions” on page 58
- “Flat file functions” on page 58
- “Java applet functions” on page 58
- “Web registry functions” on page 59
- “Persistent functions” on page 59

General purpose functions

This set of functions helps you develop Web pages by altering data or accessing system services. You can use them to send mail, process HTTP cookies, generate HTML escape codes, and get other useful information from the system.

For example, to specify that Net.Data should exit a macro if a specific condition occurs, without processing the rest of the macro, you use the DTW_EXIT function:

```
%HTML(sort_page) {  
  
<html>  
  <head>  
    <title>This is the page title</title>  
  </head>  
  <body>  
    <center>  
      <h3>This is the Main Heading</h3>  
      <!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!>  
      <! Joe Smith sees a very short page          !>  
      <!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!>  
      %IF (customer == "Joe Smith")  
    </body>  
  </html>  
  
@DTW_EXIT()  
  
%ENDIF  
  
...  
  
</body>  
</html>  
%}
```

Another useful function is the DTW_URLESCSEQ function, which replaces characters that are not allowed in a URL with their escape values. For example, if the input variable string1 equals "Guys & Dolls", DTW_URLESCSEQ assigns the output variable to the value "Guys%20%26%20Dolls".

Math functions

These functions perform mathematical operations, letting you calculate or alter numeric data. Besides standard mathematical operations, you can also perform modulus division, specify a result precision, and use scientific notation.

For example, the function DTW_POWER raises the value of its first parameter to the power of its second parameter and returns the result, as shown in the following example:

```
@DTW_POWER("2", "-3", result)
```

DTW_POWER returns ".125" in the variable result

String functions

These functions let you manipulate characters within strings. You can change a string's case, insert or delete characters, assign a string value to another variable, plus other useful functions.

For example, you can use DTW_ASSIGN to assign a value or to change the value of a variable. You can also use this function to assign a value to or to change the value of a variable. In the following example, the variable RC is assigned to zero.

```
@DTW_ASSIGN(RC, "0")
```

Other string functions include DTW_CONCAT, which concatenates strings, and DTW_INSERT, which inserts strings at a specific position, as well many other string manipulation functions.

Word functions

These functions let you manipulate words in character strings. Most of these functions work similar to string functions, but on entire words. For example, they let you count the number of words in a string, remove words, search a string for a word.

For example, use DTW_DELWORD to delete a specified number of words from a string:

```
@DTW_DELWORD("Now is the time", "2", "2", result)
```

DTW_DELWORD returns the string "Now time".

Other word functions include DTW_WORDLENGTH, which returns the number of characters in a word, and DTW_WORDPOS, which returns the position of a word within a string.

Table functions

You can use these functions to generate reports or forms using the data in a Net.Data table variable. You can also use these functions to create Net.Data tables, and to manipulate and retrieve values in those tables. Table variables contain a set of values and their associated column names. They provide a convenient way to pass groups of values to a function.

For example, DTW_TB_APPENDROW appends a row to the table. In the following example, Net.Data appends ten rows to the table, myTable:

```
@DTW_TB_APPENDROW(myTable, "10")
```

Additionally, DTW_TB_DUMP, returns the contents of a macro table variable, enclosed in <pre></pre> tags, with each row of the table displayed on a different line. And DTW_TB_CHECKBOX returns one or more HTML check box input tags from a macro table variable.

Flat file functions

Use the flat file interface (FFI) functions to open, read, and manipulate data from flat file sources (text files), as well as store data in flat files.

For example, DTWF_APPEND, writes the contents of a table variable to the end of a file, and DTWF_DELETE deletes records from a file.

Additionally, the FFI functions allow file locking with DTWF_CLOSE and DTWF_OPEN. DTWF_OPEN locks a file so that another request cannot read or update the file. DTWF_CLOSE releases the file when Net.Data is done with it, allowing other requests to access the file.

Java applet functions

Use the Java Applet functions to easily generate <applet> and <param> tags to your web page based on Net.Data variables.

For example, if you have an applet named myApplet, and you want to pass some parameters to the applet, including a table variable, you could do the following:

```
%define REMOTE_USER = %ENVVAR
%define myTable = %TABLE(all)
...
%HTML(report) {
...
@DTWA_myApplet(REMOTE_USER, myTable)
...
%}
```

This would tell Net.Data to generate an <applet> tag, as well as a <param> tag for each of the values in the table and for the value of the REMOTE_USER environment variable.

In addition, you can pass a single column of a table. For example:

```
@DTWA_myApplet(REMOTE_USER, DTW_COLUMN(mycol)myTable)
```

This example passes the mycol column of the Net.Data table variable myTable.

Web registry functions

Use the Web registry functions to maintain registries and the entries they contain. A Web registry is a file with a key maintained by `Net.Data` to allow you to add, retrieve, and delete entries easily.

For example, `DTWR_ADDENTRY` adds entries, while `DTWR_DELENTY` deletes entries. `DTWR_LISTSUB` returns information about the registry entries in an OUT table parameter, and `DTWR_UPDATEENTRY` replaces the existing values for a specified registry entry with a new value.

Persistent functions

The persistent macro functions support transaction processing in `Net.Data` by helping you define which macro blocks are persistent within a single transaction. Use these functions to define the start and end of a transaction, which HTML blocks are persistent throughout the transaction, the scope of the variables within the transaction, and whether to commit or rollback changes within the transaction.

For example, `DTW_ACCEPT` identifies the transaction handle for a transaction, while `DTW_TERMINATE` identifies the final HTML block in the transaction. `DTW_RTVHANDLE` generates a unique transaction handle for blocks in the transaction. You can use `DTW_COMMIT` and `DTW_ROLLBACK` to initiate commits and rollbacks during the transaction.

See Chapter 7, “Transaction management with persistent macros,” on page 93 for more information. Also see the built-in functions chapter in *Net.Data Reference* for a list of valid persistent functions with syntax and examples.

Generating document markup

`Net.Data` dynamically generates Web documents to be used by a client application such as a Web browser. The following sections describe the various constructs you can use to format documents with `Net.Data` macros. See the language constructs chapter in *Net.Data Reference* for the specific syntax information for each.

HTML blocks

The client application invokes `Net.Data` by specifying both the macro name and the name of one of the macro's entry points. The entry point to the macro is an HTML block. These blocks contain the `Net.Data` language statements and text presentation statements that generate the resulting page.

Because the entry point block drives the execution of the macro, at least one entry point must exist in a macro. There can be multiple HTML blocks, but only one is executed per client request. And, with each request a single document is returned to the client. To create an application consisting of many client documents, you can invoke `Net.Data` multiple times to process various HTML blocks in various macros using standard navigation techniques, such as links and forms.

Any text presentation statements can appear in an HTML block, as long as the statements are valid for the client. For example, HTML blocks can contain HTML or JavaScript. The JavaScript is not executed by `Net.Data`, but is sent along with the rest of the output to the client for execution and display. In an HTML block, you can also include function calls, variable references, and `INCLUDE` statements. The following example shows a common use of an HTML block in a `Net.Data` macro:

```
%HTML(input){
<h1>Hardware Query Form</h1>
<hr/>
<form method="post" action="report">
<dl>
</dt>What hardware do you want to list?
<dd><input type="radio" name="hardware" value="MON" checked />Monitors</dd>
<dd><input type="radio" name="hardware" value="PNT" />Pointing devices</dd>
<dd><input type="radio" name="hardware" value="PRT" />Printers</dd>
<dd><input type="radio" name="hardware" value="SCN" />Scanners</dd>
</dl>
```

```

<hr />
<input type="submit" value="Submit" />
</form>
%}

%FUNCTION(DTW_SQL) myQuery() {
SELECT MODNO, COST, DESCRIP FROM EQPTABLE WHERE TYPE='$(hardware)'
%REPORT{
<b>Here is the list you requested:</b><br />
%ROW{
<hr />
$(N1): $(V1)    $(N2): $(V2)
</p>
$(V3)
%}
%}
%}

%HTML(report){
@myQuery()
%}

```

You can invoke the Net.Data macro from an HTML link.

```

<a href="http://www.ibm.com/cgi-bin/db2www/equip1st.dtw/input">
  List of hardware</a>

```

When the application user clicks on this link, the Web browser invokes Net.Data, and Net.Data parses the macro. When Net.Data begins processing the HTML block specified on the invocation, in this case input, it begins to process the text inside the block. Anything that Net.Data does not recognize as a Net.Data macro language construct, it sends to the browser for display.

After the user makes a selection and presses the Submit button, the client requests the action specified in the action attribute of the HTML form. This action specifies a call to the output HTML block of the macro. Net.Data then processes the output HTML block, just as it did with the input HTML block.

Net.Data then processes the myQuery() function call, which in turn invokes the SQL Language Environment FUNCTION block. After replacing the \$(hardware) variable reference in the SQL statement with the value returned from the input form, Net.Data runs the query. At this point, Net.Data resumes processing the report, displaying the results of the query according to the text presentation statements specified in the REPORT block.

After Net.Data completes the REPORT block processing, it returns to the output HTML block, and finishes processing.

Report blocks

Use the REPORT block language construct to format and display data output from a FUNCTION block. This output is typically table data, although any valid combination of text, macro variable references, and function calls can be specified. A table name can optionally be specified on the REPORT block. If you do not specify a table name, Net.Data uses the table data from the first output table in the FUNCTION parameter list.

The REPORT block has three parts, each of which is optional:

- Header information, which contains text that is displayed once before the table row data.
- A ROW block, which contains text and table variables that are displayed once for each row of the result table.
- Footer information, which contains text that is displayed once after the table row data.

Example:

```

%REPORT{
<h2>Query Results</h2>
<p>Select a name for details.
<table border=1>
  <tr>
    <td>Name</td>
    <td>Location</td></tr>
  %ROW{
  <tr>
    <td>
<a href="/cgi-bin/db2www/name.dtw/details?name=$(V1)&loc=$(V2)">$(V1)</a>
    </td>
    <td>$(V2)</td>
  </tr>
  %}
</table>
%}

```

REPORT block guidelines

Use the following guidelines when creating REPORT blocks:

- To avoid displaying any table output from the ROW block, leave the ROW block empty or omit it entirely.
- Use Net.Data-provided variables inside the REPORT block to access the data in the Net.Data macro results table. These variables are described in “Table processing variables” on page 50. For additional detail, see the Report Variables section in the *Net.Data Reference*.
- To provide header and footer information, provide the text before and after the ROW block. Net.Data processes everything it finds before a ROW block as header information. Net.Data processes everything it finds after the ROW block as footer information. As with the HTML block, Net.Data treats everything in the header, ROW, and footer blocks that is not recognized as macro language constructs as text presentation statements and sends these statements to the browser.
- You can call functions and reference variables in a REPORT block.
- To have Net.Data print a default report using pre-formatted text, do not include the REPORT block in the macro. The following example shows the default report format when the function is called in an HTML block:

```

SHIPDATE | RECDATE | SHIPNO |
-----|-----|-----|
25/05/1997 | 30/05/1997 | 1495194B |
-----|-----|-----|
25/05/1997 | 28/05/1997 | 2942821G |
-----|-----|-----|

```

- To use the HTML tags instead of the pre-formatted text, set DTW_HTML_TABLE to YES.
- To disable the printing of the a default report, set DTW_DEFAULT_REPORT to NO or by specifying an empty REPORT block. For example:

```
%REPORT{%
```

Example: customizing a report

The following example shows how you can customize report formats using special variables and HTML tags. It displays the names, phone numbers, and FAX numbers from the table CustomerTbl:

```

%DEFINE SET_TOTAL_ROWS="YES"
...
%FUNCTION(DTW_SQL) custlist() {
  SELECT Name, Phone, Fax FROM CustomerTbl
  %REPORT{
<i>Phone Query Results:</i>
<br />
=====
<br />
  %ROW{
Name: <b>$(V1)</b>

```

```

<br />
Phone: $(V2)
<br />
Fax: $(V3)
<br />
-----
<br />
  %}
  Total records retrieved: $(TOTAL_ROWS)
  %}
  %}

```

The resulting report looks like this in the Web browser:

```

Phone Query Results:
=====
Name: Doen, David
Phone: 422-245-1293
Fax: 422-245-7383
-----
Name: Ramirez, Paolo
Phone: 955-768-3489
Fax: 955-768-3974
-----
Name: Wu, Jianli
Phone: 525-472-1234
Fax: 525-472-1234
-----
Total records retrieved: 3

```

Net.Data generated the report by:

1. Printing *Phone Query Results*: once at the beginning of the report. This text, along with the separator line, is the header part of the REPORT block.
2. Replacing the variables V1, V2, and V3 with their values for Name, Phone, and Fax respectively for each row as it is retrieved.
3. Printing the string *Total records retrieved*: and the value for TOTAL_ROWS once at the end of the report. (This text is the footer part of the REPORT block.)

Multiple REPORT blocks

You can specify multiple REPORT blocks within a single FUNCTION or MACRO FUNCTION block to generate multiple reports with one function call.

Typically, you would use multiple REPORT blocks with the DTW_SQL language environment with a function that calls a stored procedure, which returns multiple result sets (see “Stored procedures” on page 85). However, multiple REPORT blocks can be used with any language environment to generate multiple reports.

To use multiple REPORT blocks, place a Net.Data table variable in the function parameter list. If more result sets are returned from the stored procedure than the number of REPORT blocks you have specified, and if the Net.Data built-in function DTW_DEFAULT_REPORT = "MULTIPLE", then default reports are generated for each table that is not associated with a report block. If no report blocks are specified, and if DTW_DEFAULT_REPORT = "YES", then only one default report will be generated. Note that for the SQL language environment only, a DTW_DEFAULT_REPORT value of "YES" is equivalent to a value of "MULTIPLE".

Examples: The following examples demonstrate ways in which you can use multiple report blocks.

To display multiple reports using default report formatting:

Example 1: DTW_SQL language environment

```
%DEFINE DTW_DEFAULT_REPORT = "MULTIPLE"
%FUNCTION (dtw_sql) myStoredProc (OUT table1, table2) {
    CALL myproc  %}
```

In this example, the stored procedure myproc returns two result sets, which are placed in table1 and table2. Because no REPORT blocks are specified, default reports are displayed for both tables, table1 first, then table2.

Example 2: MACRO_FUNCTION block. In this example, two tables are passed into the MACRO_FUNCTION block. When DTW_DEFAULT_REPORT="MULTIPLE" is specified, Net.Data generates reports for both tables.

```
%DEFINE DTW_DEFAULT_REPORT = "MULTIPLE"
%MACRO_FUNCTION multReport (INOUT tablename1, tablename2) {
%}
```

In this example, two tables are passed into the MACRO_FUNCTION multReport. Again, Net.Data displays default reports for the two tables in the order in which they appear in the MACRO FUNCTION block parameter list, table1 first, then table2.

Example 3: DTW_REXX language environment

```
%DEFINE DTW_DEFAULT_REPORT = "YES"
%FUNCTION (dtw_rexx) multReport (INOUT table1, table2) {
    SAY 'Generating multiple default reports...<br />'
%}
```

In this example, two tables are passed into the REXX function multReport. Because DTW_DEFAULT_REPORT="YES" is specified, Net.Data displays a default report for the first table only.

To display multiple reports by specifying REPORT blocks for display processing:

Example 1: Named REPORT blocks

```
%FUNCTION(dtw_sql) myStoredProc (OUT table1, table2) {
    CALL myproc (table1, table2)

    %REPORT(table2) {
        ...
        %ROW { .... %}
        ...
    %}

    %REPORT(table1) {
        ...
        %row { .... %}
        ...
    %}
%}
```

In this example, REPORT blocks have been specified for both of the tables passed in the FUNCTION block parameter list. The tables are displayed in the order they are specified on the REPORT blocks, table2 first, then table1. By specifying a table name on the REPORT blocks and the CALL statement, you can control the order in which the reports are displayed.

Example 2: Unnamed REPORT blocks

```
%FUNCTION(dtw_sql) myStoredProc (OUT table1, table2) {
    CALL myproc

    %REPORT {
        ...
        %ROW { .... %}
        ...
    %}
```

```

%}
%REPORT {
    ...
    %ROW { .... %}
    ...
%}
%}

```

In this example, REPORT blocks have been specified for two result sets returned from myproc. Because there are no table names specified on the REPORT blocks, the REPORT blocks are executed for the first two result sets in the order in which they are returned from the stored procedure.

To display multiple reports using a combination of default reports and REPORT blocks:

Example: A combination of default reports and REPORT blocks

```

%DEFINE DTW_DEFAULT_REPORT = "MULTIPLE"
%FUNCTION(dtw_system) editTables (INOUT table1, table2, table3) {
    %EXEC{ /qsys.lib/mylib.lib/mypgm.pgm %}
    %REPORT(table2) {
        ...
        %ROW { .... %}
        ...
    %}
%}

```

In this example, only one REPORT block is specified, and because it specifies a table name of table2, it uses this table to display its report. Because there are fewer REPORT blocks specified the number of result sets returned from the stored procedure, default reports are displayed for the remaining for the remaining result sets: first, a default report for table1; then a default report for table3.

Guidelines and restrictions for multiple REPORT blocks: Use the following guidelines and restrictions when specifying multiple REPORT blocks in a FUNCTION or MACRO_FUNCTION block.

Guidelines:

- You can specify one REPORT block per result set.
- Specify REPORT blocks for multiple tables in the order in which you want them to be processed.
- To specify default processing when there is not a REPORT block specified for a table, define DTW_DEFAULT_REPORT = "MULTIPLE". When Net.Data builds the Web page, it displays default reports for tables after it displays the reports for tables having REPORT blocks. Note that setting DTW_DEFAULT_REPORT = "YES" will result in the generation of a default report for one table only, when a REPORT block has not been specified. The exception is in the SQL language environment, where a value of YES will result in the same processing as MULTIPLE.
- To prevent Net.Data from displaying tables that do not have REPORT blocks, set DTW_DEFAULT_REPORT = "NO".
- When using the DTW_SAVE_TABLE_IN variable with a function that returns more than one result set, the first result set returned from the function is assigned to the DTW_SAVE_TABLE_IN table.
- Multiple report blocks can be used with any language environment.

Restrictions:

- The values of all report variables in a function, such as START_ROW_NUM and RPT_MAX_ROWS, apply to all the REPORT blocks in that function. You cannot modify the value of a report variable for individual REPORT blocks.
- The MESSAGE block must be located either before or after a list of REPORT blocks, and not between REPORT blocks.
- If the first report block specifies a table name, then all report blocks must specify table names.

- If the first report block does not specify a table name, then none of the report blocks can specify table names.

Conditional logic and looping in a macro

Net.Data lets you incorporate conditional logic and looping in your Net.Data macros using the IF and WHILE blocks.

IF and WHILE blocks use a condition list that helps you test one or more conditions, and then to perform a block of statements based on the outcome of the condition test. The condition list contains logical operators, such as = and <=, and terms, which are made up of quoted strings, variables, variable references, and function calls. Quoted strings can contain variable references and functions calls, as well. You can nest the condition list.

The following sections describe conditional logic and looping:

- “Conditional logic: IF blocks”
- “Looping constructs: WHILE blocks” on page 67

Conditional logic: IF blocks

Use the IF block for conditional processing in a Net.Data macro. The IF block is similar to IF statements in most high-level languages because it provides the ability to test one or more conditions, and then to perform a block of statements based on the outcome of the condition test.

You can specify IF blocks almost anywhere in a macro and can nest them. The syntax of an IF block is shown in the language constructs chapter in *Net.Data Reference*.

IF Block Rules: The rules for IF block syntax are determined by the block's position in the macro. The elements allowed in the executable block of statements of an IF block depend on the location of the IF block itself.

- Any element that is valid in the block containing the IF block is valid within that IF block. For example, if you specify an IF block inside an HTML block, any element that is allowed in the HTML block is allowed in the IF block, such as INCLUDE statements and WHILE blocks.

```
%HTML block
...
  %IF block
  ...
    %INCLUDE
  ...
    %WHILE
  ...
  %ENDIF
%}
```

- Similarly, if you specify the IF block outside of any other block in the declaration part of the Net.Data macro, only those elements allowed outside of any other block (such as a DEFINE block or FUNCTION block) are allowed in the IF block.

```
%IF
...
  %DEFINE
...
  %FUNCTION
...
%ENDIF
```

- When an IF block is nested within an IF block that is outside of any other block in the declaration part, it can use any element that the outside block can use. When an IF block is nested within another block that is in an IF block, it takes on the syntax rules for the block it is inside.

For example, a nested IF block must follow the rules used when it is inside an HTML block.

```

%IF
...
%HTML {
...
%IF
...
%ENDIF
%}
...
%ENDIF

```

Exception: Do not specify a ROW block in an IF block.

IF Block String Comparison

Net.Data processes the IF block condition list in one of two ways based on the contents of the terms making up the conditions. The default action is to treat all terms as strings, and to perform string comparisons as specified in the conditions. However, if the comparison is between two strings representing integers, then the comparison is numeric. Net.Data assumes a string is numeric if it contains only digits, optionally preceded by a '+' or '-' character. The string cannot contain any non-digit characters other than the '+' or '-'. Net.Data does not support numerical comparison of non-integer numbers.

Examples of valid integer strings:

```

+1234567890
-47
000812
92000

```

Examples of invalid integer strings:

```

- 20      (contains blank characters)
234,000   (contains a comma)
57.987    (contains a decimal point)

```

Net.Data evaluates the IF condition at the time it executes the block, which can be different than the time it is originally read by Net.Data. For example, if you specify an IF block in a REPORT block, Net.Data does not evaluate the condition list associated with the IF block when it reads the FUNCTION block definition containing the REPORT block, but rather when it calls the function and executes it. This is true for both the condition list part of the IF block and the block of statements to be executed.

IF Block Example: A macro containing IF blocks inside other blocks

```

%{ This macro is called from another macro, passing the operating system
and version variables in the form data.
%}

%IF (platform == "AS400")
%IF (version == "V3R2")
%INCLUDE "as400v3r2_def.hti"
%ELIF (version == "V3R7")
%INCLUDE "as400v3r7_def.hti"
%ELIF (version == "V4R1")
%INCLUDE "as400v4r1_def.hti"
%ENDIF
%ELSE
%INCLUDE "default_def.hti"
%ENDIF

%MACRO_FUNCTION numericCompare(IN term1, term2, OUT result) {
%IF (term1 < term2)
@dtw_assign(result, "-1")
%ELIF (term1 > term2)
@dtw_assign(result, "1")
%ELSE
@dtw_assign(result, "0")
%ENDIF
}

```

```

%}
%HTML(report){
  %WHILE (a < "10") {
    outer while loop #$(a)<br />
    %IF (@dtw_rdivrem(a,"2") == "0")
      this is an even number loop<br />
    %ENDIF
    @DTW_ADD(a, "1", a)
  %}
%}

```

Looping constructs: WHILE blocks

Use the WHILE block to perform looping in a Net.Data macro. Like the IF block, the WHILE block provides the ability to test one or more conditions, and then to perform a block of statements based on the outcome of the condition test. Unlike the IF block, the block of statements can be executed any number of times based on the outcome of the condition test.

You can specify WHILE blocks inside HTML blocks, REPORT blocks, ROW blocks, MACRO_FUNCTION blocks, and IF blocks, and you can nest them. The syntax of a WHILE block is shown in the language constructs chapter of *Net.Data Reference*.

Net.Data processes the WHILE block exactly the same way it processes the IF block, but re-evaluates the condition after each execution of the block. And, like any conditional looping construct, it is possible for processing to go into an infinite loop if the condition is coded incorrectly.

Example: A macro with a WHILE block

```

%DEFINE loopCounter = "1"

%HTML(build_table) {
  %WHILE (loopCounter <= "100") {
    %{ generate table tag and column headings %}
    %IF (loopCounter == "1")
      <table border>
      <tr>
      <th>Item #
      <th>Description
    %ENDIF

    %{ generate individual rows %}
    <tr>
    <td>$(loopCounter)
    <td>@getDescription(loopCounter)

    %{ generate end table tag %}
    %IF (loopCounter == "100")
    %ENDIF

    %{ increment loop counter %}
    @DTW_ADD(loopCounter, "1", loopCounter)
  %}
%}

```


Chapter 6. Using language environments

Net.Data supplies language environments that you use to access data sources and to execute application programs containing business logic. For example, the SQL language environment lets you pass SQL statements to a DB2 database, and the REXX language environment lets you invoke REXX programs. You can also use the SYSTEM language environment to execute a program or issue a command.

With Net.Data, you can add user-written language environments in a pluggable fashion. Each user-written language environment must support a standard set of interfaces that are defined by Net.Data and must be implemented as a service program. For complete details on how to create a user-written language environment, see the *Net.Data Language Environment Interface Reference*.

Figure 7 shows the relationship between the Web server, Net.Data, and the Net.Data language environments.

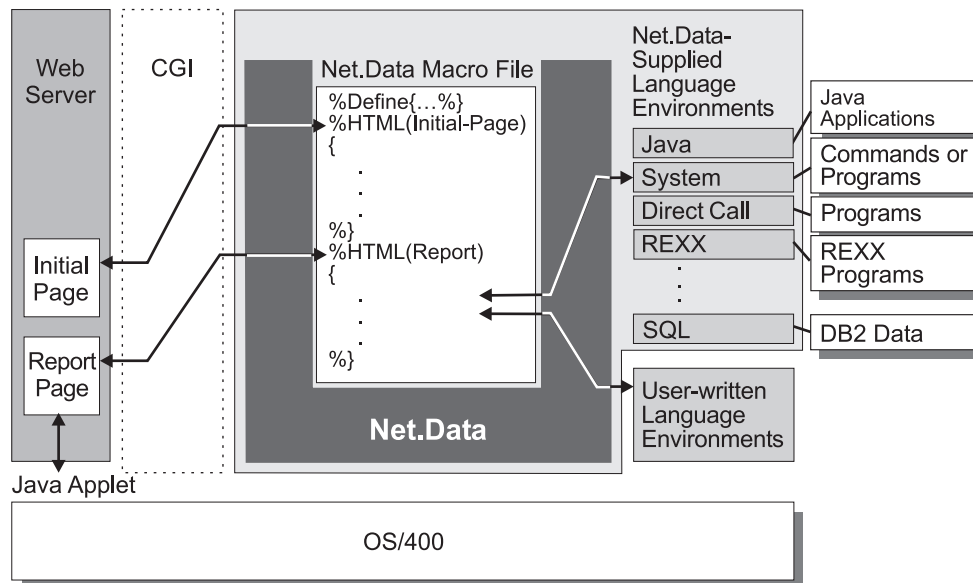


Figure 7. The Net.Data Language Environments

The following sections describe the Net.Data language environments and how to use them in your macros:

- “Overview of Net.Data-supplied language environments” on page 70
- “Calling a language environment” on page 70
- “Direct Call language environment” on page 72
- “Java Application language environment” on page 76
- “REXX language environment” on page 78
- “SQL language environment” on page 82
- “System language environment” on page 91

For information about improving performance when using the language environments, see “Optimizing the language environments” on page 103.

Overview of Net.Data-supplied language environments

Net.Data provides language environments that let you access data and programming resources for your application.

Table 3 provides a brief description of each language environment.

Table 3. *Net.Data Language Environments*

Language Environment	Environment Name	Description
Direct Call	DTW_DIRECTCALL	The Direct Call language environment supports calls to external programs that are written using a high-level programming language such as RPG, COBOL, and C/C++.
Java Application	DTW_JAVAPPS	Net.Data supports your existing Java applications with the Java language environment.
REXX	DTW_REXX	The REXX language environment interprets internal REXX programs that are specified in a FUNCTION block of the Net.Data macro, or it can execute external REXX programs stored in a separate file.
SQL	DTW_SQL	The SQL language environment executes SQL statements through DB2. The results of the SQL statement can be returned in a table variable. The results of the ODBC statement can be returned in a table variable.
System	DTW_SYSTEM	The System language environment supports executing commands and calling external programs.

Calling a language environment

To call a language environment:

- Use a FUNCTION statement to define a function that calls the language environment by supplying language statements or an %EXEC statement.
- Use a function call to the language environment.

For example:

```
%FUNCTION(DTW_SQL) custinfo() {  
  select customer, custno from customer.data  
  %}  
...  
%HTML(REPORT) {  
  @custinfo()  
  %}
```

Guidelines for handling error conditions

When an error is detected in a language environment function, the language environment sets the Net.Data RETURN_CODE variable with an error code.

You can use the following resources to handle error conditions:

- The Net.Data-supplied language environments return error codes that are documented in *Net.Data Messages and Codes Reference*.
- The database language environments, such as the SQL language environment set the RETURN_CODE variable to the SQLCODE returned by the database, and the SQL_STATE variable to the SQLSTATE returned by the database. See the messages and codes documentation for your DBMS to learn more about the SQLCODEs and SQLSTATEs used by your DBMS.

Security

Ensure that the user ID Net.Data is running under has the proper authority to access any object that may be referenced by a language environment statement. For example, SQL language environment executes SQL statements, so that the user ID under which Net.Data executes must have the authority to access the database resources, in order to execute successfully.

Direct Call language environment

The Direct Call language environment allows you to call programs that are written in a high-level language such as C, RPG, COBOL, and CL. Parameters can be passed *to* the program, and parameter values can be received *from* the program, enabling easy integration of existing programs with Net.Data and allowing users to use existing programming skills to code complicated business logic.

Calling programs

To call a program, define a function that uses the Direct Call (DTW_DIRECTCALL) language environment and that includes a path to the program that is to be called in an EXEC statement. For example:

```
%function(DTW_DIRECTCALL) dc1() {  
    %EXEC { /QSYS.LIB/NETDATA.LIB/MYPGM.PGM %}  
%}
```

You can shorten the path to the program if you use the EXEC_PATH configuration variable to define paths to directories that contain programs. See “EXEC_PATH” on page 16 to learn how to define the EXEC_PATH configuration variable.

Supported Language Environment Variables

The Direct Call language environment supports the DTW_PAD_PGM_PARMS variable, which indicates whether parameters that are to be passed to a program are to be padded with blanks up to the precision specified. See *Net.Data Reference* for description, syntax, and examples for this variable. See “Passing parameters to programs” for more information on passing parameters to programs.

Passing parameters to programs

Pass parameters to a program by specifying on the function definition the data type of the parameter and whether the parameter to be passed is an input-only (IN), output-only (OUT), or input/output (INOUT) parameter. For example:

```
%function(DTW_DIRECTCALL) dc2(IN CHAR(3) p1,  
                               INOUT INTEGER p2,  
                               OUT DECIMAL(7,2) p3) {  
    %EXEC { /QSYS.LIB/NETDATA.LIB/MYPGM.PGM %}  
%}
```

In the above example, the Direct Call language environment passes three parameters, a character variable, an integer, and a packed decimal variable to the program MYPGM. You can pass up to 50 parameters to the called program. Only parameters specified with data types are passed to the program. The Direct Call language environment converts the string corresponding to the parameter to the internal representation of the data type. The language environment then passes pointers to the internal representation of the variables to the called program, in the order specified on the function definition.

Because pointers to the variables are passed to the program, the program can change the value of the variable. However, only OUT or INOUT variables that are changed by the program are reflected back in the macro that called the language environment.

Supported data types

Table 4 on page 73 lists the data types that are supported by the Direct Call language environment. Not all of the data types are supported by each high-level language.

Table 4. Direct Call Data Types

Data Type	Usage Notes
CHAR(<i>n</i>) CHARACTER(<i>n</i>) CHARACTER	A character string. If <i>n</i> is specified, it must be greater than zero. If the string is not specified, it is assumed to be one character. Because all strings passed from the Direct Call language environment are null-terminated, the language environment allocates <i>n</i> +1 bytes (1 byte for the NULL terminator). Strings that exceed <i>n</i> are truncated.
VARCHAR(<i>n</i>)	A variable-length character string, where <i>n</i> is greater than zero, and less than or equal to 32740. The string is null-terminated and the language environment allocates <i>n</i> +2+1 bytes (2 bytes to store the string length, 1 byte for the null-terminator). Strings that exceed <i>n</i> are truncated. The first two bytes of the string contain the string length (binary value). If the parameter is defined as OUT (output only), the string length is set to zero before the variable is passed to the called program.
INTEGER INT	A signed binary integer, 4 bytes long.
SMALLINT	A signed binary integer, 2 bytes long
FLOAT(<i>p,s</i>)	A single-precision or double-precision, floating point number. For single precision, <i>p</i> must be greater than 0 and less than 25. For double precision, <i>p</i> must be greater than 24 and less than 54. The precision (<i>p</i>) and scale (<i>s</i>) are only used when converting data to a displayable format; for example, to a string.
REAL(<i>p,s</i>)	A single-precision floating point number. <i>p</i> must be greater than 0 and less than 25. The precision (<i>p</i>) and scale (<i>s</i>) are only used when converting data to a displayable format; for example, to a string.
DOUBLE (<i>p,s</i>) DOUBLEPRECISION(<i>p,s</i>)	A double-precision floating point number. <i>p</i> must be greater than 0 and less than 53. The precision (<i>p</i>) and scale (<i>s</i>) are only used when converting data to a displayable format; for example, to a string.
NUMERIC(<i>p,s</i>)	A zoned decimal number, with precision <i>p</i> and scale <i>s</i> . The value of <i>p</i> must be greater than 0 and less than 32.
DEC(<i>p,s</i>) DECIMAL(<i>p,s</i>)	A packed decimal number, with precision <i>p</i> and scale <i>s</i> . The value of <i>p</i> must be greater than 0 and less than 32.
DTWTABLE	A special data type used to pass a Net.Data table to the called program. The Direct Call language environment passes a pointer to the table, which can be manipulated using the Net.Data language environment interface table functions.

Parameters that are defined to be numeric can include the currency symbol and three-digit separators. The Direct Call language environment removes the currency symbol and three-digit separators when converting a numeric variable from string form to its internal form, before passing the variable to the program. Net.Data retrieves the currency symbol, decimal format, and three-digit separator characters from the process attributes of the process in which Net.Data is running.

Null-terminated string parameters

If DTW_PAD_PGM_PARMS is set to NO in the configuration file or within the macro, the Direct Call language environment passes string values to your program using a null terminator character (value x'00'). This requires you to write code to handle the string (unless you are using C or C++, which expect null terminated strings).

For example, if you define the parameter field as CHAR(10), but pass a string value that is 5 bytes long, Net.Data puts the null terminator after the fifth byte. Passing the value "12345" as a string in a CHAR(10) field yields:

```
x'F1F2F3F4F500.....'
```

The bytes following the null terminator are undefined (you cannot assume that the bytes are null or blank).

Because the string is null terminated and contains uninitialized bytes after the null terminator, you cannot use the string in an RPG or COBOL program. For example, if you use the string in a comparison operation, the operation does not yield valid results. The program does not expect the string to contain the null terminator and expects the string to be padded with blanks at the end.

You can use string handling functions within your program to extract the string value or use the VARCHAR data type. This method gives the length of the string in the first two bytes.

If DTW_PAD_PGM_PARMS is set to YES in the configuration file or within the macro, the Direct Call language environment passes string values to your program with the values padded to the right with blanks up to the precision length. Using the same example as above, but with DTW_PAD_PGM_PARMS set to YES, passing the value "12345" as a string in a CHAR(10) field yields:

```
x'F1F2F3F4F54040404000'
```

Because the string has a length of 5, which is less than the specified precision, blanks are inserted after the value up to the precision length. Programs written in languages such as RPG can now use the parameter without the need to handle NULL-terminated strings.

Common Errors when Passing Parameters

The following list describes errors that can occur when calling programs and passing parameters to the program using the Direct Call language environment. Tips for avoiding these errors are provided.

Parameter Mismatch Errors

Ensure that the number and order of parameters match the number and order in which they appear in the parameter list of the called program.

Data Type Errors

Ensure that the data type specified for a parameter matches the data type expected by the called program. There might be data types supported by the Direct Call language environment that are not supported by the high level programming language used to create the called program.

Length Errors

Ensure that the lengths defined for parameters are correct and match the lengths specified in the called program. Specifying a length that is shorter than the declared length of the called program might corrupt storage and cause Net.Data not to function properly.

Returning values from programs

Some high level programming languages, such as C, can return an integer on the program call. The integer can be retrieved by specifying the RETURNS keyword in the function definition, For example:

```
%function(DTW_DIRECTCALL) dc3(IN CHAR(3) p1) RETURNS(retval) {  
  %EXEC { /QSYS.LIB/NETDATA.LIB/MYPGM.PGM %}  
%}
```

When the function call completes successfully, the parameter *retval* contains the value returned by the program.

Direct Call language environment example

In this example, the macro calls a program and passes several parameters. The source for the program follows the macro, and is written in RPG and CL. The program that is called accepts two integer parameters. It copies the first parameter (the input parameter) to the second parameter (the output parameter).

Macro:

```

%define ilepgm = "/QSYS.LIB/NETDATADEV.LIB/TDCCLI01.PGM"
%define out1 = "0"

%FUNCTION(DTW_DIRECTCALL) dcFunction(IN INT inp1,
                                     OUT INT outp2)
{ %EXEC { $(ilepgm) %} %}

%HTML(REPORT) {
@dcFunction("123", out1)
The value of out1 is: "$(out1)"
%}

```

ILE RPG program:

```

DINP1          S           10I00
DOUTP2         S           10I00
C*
C   *ENTRY      PLIST
C               PARM           INP1
C               PARM           OUTP2
C*
C               Z-ADD      INP1      OUTP2
C*
C               SETON
C                                     LR

```

CL program:

```

PGM PARM(&INP1; &OUTP2;)

DCL VAR(&INP1;) TYPE(*CHAR) LEN(4)
DCL VAR(&OUTP2;) TYPE(*CHAR) LEN(4)

CHGVAR VAR(&OUTP2;) VALUE(&INP1;)
ENDPGM

```

Java Application language environment

The Java Application language environment allows you to call Java programs, enabling easy integration of Java applications with Net.Data. The Java Application language environment was first introduced in OS/400 V4R4.

To use the Java Application language environment, complete the configuration steps documented in "Setting up the Java application language environment" on page 20.

Calling Java programs

To call a Java program, define a function that uses the Java Application (DTW_JAVAPPS) language environment. Specifying a function name that represents the class name of the Java program.

Example: Calls a Java program `helloWorld.java`:

```
%function(DTW_JAVAPPS) helloWorld() { %}
```

The Java Application language environment expects Java programs to contain a method identifier for 'main,' the first method that is run in a Java program. When the language environment invokes an application, the application has access to stdin and stdout. There is no form data in stdin because Net.Data has already read the data.

Important: Before calling Java applications, set the `DTW_JAVA_CLASSPATH` path configuration variable so that the Java class can be found. See "DTW_JAVA_CLASSPATH" on page 16 for the syntax of this variable.

Passing parameters to Java programs

Pass parameters to a Java program by specifying the parameters to be passed on the function definition. Specify only string parameters that are input-only (IN), or input or output (INOUT).

Example: The IN parameter `p1` is to be passed on the function call

```
%function(DTW_JAVAPPS) jv1(IN p1) { %}
```

The Java Application language environment does not support Java programs that update variables because it cannot pass the updated values back to the macro.

Java Application language environment example

In this example, the Net.Data macro calls a Java program, `echoString`. The macro passes two string parameters to the Java language environment. The first string tells the Java program whether to use italic or bold highlighting for the second parameter, a text string, before printing the second parameter to standard output (stdout). Because the program passes "I", for italics, the Web server displays the text string *Hello World*, in italics, at the browser. The source for the Java program follows the macro.

Macro:

```
%FUNCTION(DTW_JAVAPPS) echoString(textAttribute, text){ %}

%HTML(runjava){
@echoString("I","Hello World")
%}
```

Java program:

```
class echoString {
    public static void main (String args[]) {
        if (args[0].equals("I"))
            System.out.println("<i>" + args[1] + "</i>");
    }
}
```

```
    else
        System.out.println("<b>" + args[1] + "</b>");
}
```

REXX language environment

The REXX language environment allows you to run REXX programs written to run in the DTW_REXX environment. The Net.Data REXX language environment provides controls that allow REXX programs to easily return large amounts of data.

Net.Data also provides support for the REXX SAY statement that directs the output to the client.

Support for Variables: To allow REXX programs to easily return large amounts of data, Net.Data automatically adds code to the beginning of the REXX program and appends code to the end of the REXX program. This code is designed to manipulate variables that were provided on the DTW_REXX function statement.

Use of REXX subroutines and functions: Since Net.Data adds code to the front of the REXX program and appends code to the end of the REXX program, the main REXX routine must end with the last statement of the REXX program. If you use REXX subroutines or functions you must insure that the last statement of the REXX program is associated with the main REXX routine. The following is an example of using a subroutine and function in a REXX program that is written to run in the DTW_REXX environment:

```
%function(DTW_REXX) genData(out s1,s2) {
    call subrtn1
    s2=funrtn1()
    signal rexxEnd /* Go to end of Program */
    subrtn1: PROCEDURE EXPOSE s1
        string1 = "data for s1"
        return 0
    funrtn1: PROCEDURE
        retvar = "data for s2"
        return retvar
    rexxEnd:      /* End of Main Program */
        return 0
}%
%HTML (Report) {

    @genData(a,c)

    Value for s1: $(a)

    Value for s2: $(c)
}%
```

Use of REXX EXIT and RETURN statements: Net.Data automatically appends code to REXX programs that provide values for output variables and directs output from SAY statements to the browser. If the REXX program issues a RETURN from the main routine or issues an EXIT statement anywhere but the last statement of the REXX program, the code that was appended by Net.Data to the REXX program will not be executed. This results in the lost of output variables and output from SAY statements. If you must exit a REXX program before reaching the last statement, you should branch to the last statement in the REXX program that normally exits. If you use the RETURN or EXIT statement to end the main REXX program, it must be the last statement in the REXX program. This includes REXX comment statements. For example:

```
%function(DTW_REXX) genData(out s1,s2) {
    .....
    If S2 < 0 Then signal rexxEnd
    .....
    .....
    rexxEnd:
    /* This comment must be before the following
    RETURN statement */
    return 0
}%
```

```
%HTML (Report) {
@genData(a,c)
.....
%}
```

Invoking external REXX programs from a DTW_REXX function: You can invoke a REXX program from a DTW_REXX function using the Net.Data %EXEC statement or from a REXX program using methods provided by REXX.

When invoking an external REXX program using the Net.Data %EXEC statement, Net.Data automatically adds code to the beginning of the REXX program and appends code to the end of the REXX program to handle Output variables and direct output from REXX SAY statements to the browser.

When you use methods provided by REXX to invoke a REXX program, Net.Data does not receive control and doesn't add code to the REXX program. The REXX program being invoked must pass output back to the calling REXX program using standard REXX conventions.

Executing REXX programs

With the REXX language environment you can execute both in-line REXX programs or external REXX programs. An in-line REXX program is a REXX program that has the source of the REXX program in the macro. An external REXX program has the source of the REXX program in an external file.

To execute an in-line REXX program:

Define a function that uses the REXX (DTW_REXX) language environment and contains the REXX code in the function.

Example: A function that contains a in-line REXX program

```
%function(DTW_REXX) helloWorld() {
    SAY 'Hello World'
%}
```

To run an external REXX program:

Define a function that uses the REXX (DTW_REXX) language environment and includes a path to the REXX program that is to be run in an EXEC statement.

Example: A function that contains an EXEC statement pointing to a the external program

```
%function(DTW_REXX) externalHelloWorld() {
%EXEC{ /QSYS.LIB/REXX.LIB/REXXSRC.FILE/HELLOWORLD.MBR%}
%}
```

You can shorten the path to the program if you use the EXEC_PATH configuration variable to define paths to directories that contain programs. See “EXEC_PATH” on page 16 to learn how to define the EXEC_PATH configuration variable.

Restriction: If you are running OS/400 V3R2 or V3R7 and a REXX program uses the SAY REXX instruction to write to stdout, then insert 12 blanks at the start of the string. For example:

```
SAY '          STARTOFDATA'
```

The 12 blanks are ignored, but if they are not inserted, unpredictable results might occur.

Passing parameters to REXX programs

There are two ways to pass information to a REXX program that is invoked by the REXX (DTW_REXX) language environment, directly and indirectly.

Directly

Pass parameters directly to an external REXX program using the %EXEC statement. For example:

```
%FUNCTION(DTW_REXX) rexx1() {  
    %EXEC{/QSYS.LIB/NETDATA.LIB/QREXSSRC.FILE/CALL1.MBR $(INPARAM1) %}  
%}
```

The Net.Data variable INPARAM1 is referenced and passed to the external REXX program. The REXX program can reference the variable by using REXX PARSE ARG instruction. The parameters that are passed to the REXX program using this method are considered input parameters, and any modification to the values are not reflected back to Net.Data. (the parameters passed to the program can be used and manipulated by the program, but changes to the parameters are not reflected back to Net.Data).

Indirectly

Pass parameters indirectly, by way of the REXX program *variable pool*. When a REXX program is started, a space which contains information about all variables is created and maintained by the REXX interpreter. This space is called the variable pool.

When a REXX language environment (DTW_REXX) function is called, any function parameters that are input (IN) or input/output (INOUT) are stored in the by the REXX language environment prior to executing the REXX program. When the REXX program is invoked, it can access these variables directly. Upon the successful completion of the REXX program, the DTW_REXX language environment determines whether there are any output (OUT) or INOUT function parameters. If so, the language environment retrieves the value corresponding to the function parameter from the variable pool and updates the function parameter value with the new value. When Net.Data receives control, it updates all OUT or INOUT parameters with the new values obtained from the REXX language environment. For example:

```
%DEFINE a = "3"  
%DEFINE b = "0"  
%FUNCTION(DTW_REXX) double_func(IN inp1, OUT outp1){  
    outp1 = 2*inp1  
%}  
  
%HTML (Report) {  
    Value of b is $(b), @double_func(a, b) Value of b is $(b)  
%}
```

In the above example, the call *@double_func* passes two parameters, *a* and *b*. The REXX function *double_func* doubles the first parameter and stores the result in the second parameter. When Net.Data invokes the macro, *b* has a value of 6.

You can pass Net.Data tables to a REXX program. A REXX program accesses the values of a Net.Data macro table parameter as REXX stem variables. To a REXX program, the column headings and field values are contained in variables identified with the table name and column number. For example, in the table *myTable*, the column headings are *myTable_V.j*, and the field values are *myTable_V.i.j*, where *i* is the row number and *j* is the column number. The number of rows in the table is *myTable_ROWS* and the number of columns in the table is *myTable_COLS*.

REXX language environment example

The following example shows a macro that calls a REXX function to generate a Net.Data table that has two columns and three rows. Following the call to the REXX function, a built-in function, DTW_TB_TABLE(), is called to generate an HTML table that is sent back to the browser.

```
%DEFINE myTable = %TABLE  
%DEFINE DTW_DEFAULT_REPORT = "NO"  
  
%FUNCTION(DTW_REXX) genTable(out out_table) {  
    out_table_ROWS = 3  
    out_table_COLS = 2
```

```

/* Set Column Headings */
do j=1 to out_table_COLS
  out_table_N.j = 'COL'j
end

/* Set the fields in the row */
do i = 1 to out_table_ROWS
  do j = 1 to out_table_COLS
    out_table_V.i.j = '[' i j ']'
  end
end
end
%}

%HTML (Report) {
  @genTable(myTable)
  @DTW_TB_TABLE(myTable)
%}

```

Results:

```

<table>
  <tr>
    <th>COL1</th>
    <th>COL2</th>
  </tr>
  <tr>
    <td>[1 1]</td>
    <td>[1 2],</td>
  </tr>
  <tr>
    <td>[2 1]</td>
    <td>[2 2],</td>
  </tr>
  <tr>
    <td>[3 1]</td>
    <td>[3 2],</td>
  </tr>
</table>

```

SQL language environment

The SQL language environment allows you to run SQL statements by sending the SQL statements to a database management system (DBMS).

To use the SQL language environment, ensure you follow the configuration steps documented in “Setting up Net.Data language environments” on page 20

Executing SQL statements

You can execute any SQL statement that is supported by dynamic SQL.

To execute SQL statements, define a function that uses the SQL (DTW_SQL) language environment and contains the SQL statements in the language environment executable section of the function.

Example: An SQL function that runs an SQL SELECT statement:

```
%function(DTW_SQL) getOrders() {  
    SELECT cust, custid, custorder FROM mylibrary.customers  
%}
```

Commitment Control

The SQL language environment by default runs under commitment control and follows all rules for governing commitment control.

- Journal all files or tables that are accessed through the DTW_SQL, except when the SQL statement is SELECT.
- Optionally change the commitment level by specifying DTW_SQL_ISOLATION in the Net.Data initialization file. See “DTW_SQL_ISOLATION” on page 13 for details about the isolation levels that the SQL language environment supports.

For more information on transaction management, see “Managing transactions in a Net.Data application” on page 84.

OUT and INOUT Tables

If you specify OUT or INOUT Net.Data tables on the function definition, and the SQL statement returns result sets, the SQL language environment stores each result set in the specified tables. You can then use the table later in the macro. If an OUT table is not specified, the SQL language environment uses a default table.

Nested SQL Statements

You can call SQL functions within another SQL function. If tables are passed, then ensure you use unique table names in each of the functions; otherwise, unpredictable results might occur.

Example: Calls an SQL function from the ROW block of another SQL function

```
%define mytable1 = %TABLE  
%define mytable2 = %TABLE  
  
%FUNCTION(DTW_SQL) sql2 (IN p1, OUT t2) {  
    select * from NETDATA.STAFFINF where projno='$(p1)'  
    %REPORT {  
        %ROW { $(N1) is $(V1) %}  
    %}  
%}  
  
%FUNCTION(DTW_SQL) sql1 (OUT t1) {  
    select * from NETDATA.STAFFINF
```

```

%REPORT {
  %ROW { @sql2(V1, mytable2) %}
%}
%}

%HTML(netcall1) { @sql1(mytable1) %}

```

Supported language environment variables

The SQL language environment supports variables designed to support DB2. For example, the DATABASE variable specifies the data source that the SQL language environment connects to when executing an SQL statement. The following list specifies which variables are supported for the SQL language environment. See *Net.Data Reference* for description, syntax, and examples for these variables.

- DATABASE
- DB_CASE
- DTW_EDIT_CODES
- DTW_PAD_PGM_PARMS
- DTW_SET_TOTAL_ROWS
- LOGIN
- NULL_RPT_FIELD
- PASSWORD
- SHOWSQL
- SQL_STATE
- TRANSACTION_SCOPE

Supported data types

The SQL language environment supports the data types listed in Table 5

Table 5. Data Types

CHAR	DOUBLE	SMALLINT
DATE	DOUBLEPRECISION	TIME
DECIMAL	FLOAT	TIMESTAMP
	GRAPHIC	VARCHAR
	INTEGER	VARGRAPHIC
	REAL	

SQL language environment restrictions

Consider the following restrictions when planning your environment:

- Do not use the SQL language environment if at least one of the following conditions exists:
 - A user-defined language environment is created that uses the database access class library or the SQL call level interface (CLI) and the user-defined language environment is referenced in a macro
 - An application that uses the SQL CLI will be running in the same process as Net.Data
- You can use up to 50 local or remote database connections. When using multiple connection, consider the following restrictions:
 - Net.Data does not allow concurrent connections to the same remote database.
 - You cannot change login IDs after you have accessed a remote database if TRANSACTION_SCOPE=MULTIPLE, which is the default. See “Managing transactions in a Net.Data application” on page 84.

See “Managing multiple database connections” on page 84 for more information about these restrictions.

Managing transactions in a Net.Data application

When you modify the content of a database using insert, delete, or update statements, the modifications do not become persistent until the database receives a commit statement from Net.Data. If an error occurs, Net.Data sends a rollback statement to the database, reversing all modifications since the last commit.

The way in which Net.Data sends the commit and possible rollback statements depend on the setting of TRANSACTION_SCOPE and whether commit statements are explicitly specified in the macro. The values for TRANSACTION_SCOPE are MULTIPLE and SINGLE. The default is MULTIPLE. To set TRANSACTION_SCOPE to SINGLE, use a %DEFINE statement or a call to @DTW_ASSIGN().

SINGLE

Specifies that Net.Data issues a commit statement after each successful SQL statement. If the SQL statement returns an error, a rollback statement is issued. SINGLE transaction scope secures a database modification immediately; however, with this scope, it is not possible to undo a modification using a rollback statement later.

MULTIPLE

Specifies that Net.Data will execute all SQL statements before a commit statement is issued. Net.Data sends the commit at the end of the request, and if each SQL statement is issued successfully, the commit makes all modifications in the database persistent. If any of the statements returns an error, Net.Data issues a rollback statement at the point of the error, which sets the database back to its prior state.

By leaving TRANSACTION_SCOPE set to MULTIPLE and issuing commit statements at the end of those groups of statements that you feel qualify as a transaction, you the application developer maintain full control over the commit and rollback behavior in your application.

To issue an SQL commit statement, you can define a function that you can call in at any point in your HTML block:

```
%FUNCTION(DTW_SQL) user_commit() {
    commit
}%

...

%HTML {
    ...
    @user_commit()
    ...
}%
```

Restrictions:

The setting of TRANSACTION_SCOPE cannot be changed after a connection to the database is made. Therefore, all SQL transactions in a macro are subject to the same processing.

If you are using Net.Data as part of Net.Commerce, note that Net.Commerce has its own transaction handling and disables the transaction handling of Net.Data.

Managing multiple database connections

You can connect to up to 50 local or remote databases at a time. The SQL language environment keeps the connections active for the life of the Web server process job that Net.Data is running under. Keeping the connections active provides fast database access after the initial connection to the database. You can prevent errors by taking the following issues into consideration:

- Net.Data does not allow concurrent connections to the same remote database. If a connection exists to a remote database using one user ID (the LOGIN SQL language environment parameter) and another request is made to connect to the same remote database using a second user ID, the SQL language environment must first disconnect the existing connection, do a commit (if commitment control is being used) and then reestablish the connection using the 'new' user ID and password. The commit is required because if the connection is broken, there is no way that a rollback can be accomplished in case of an error later in the macro.
- You can change login IDs after you've accessed a remote database, if TRANSACTION_SCOPE=SINGLE. The SQL language environment disconnects the existing connection, does a commit, and reestablishes the connection using the new user ID and password.
- Do not change login IDs after you have accessed a remote database if TRANSACTION_SCOPE=MULTIPLE, which is the default. The SQL language environment automatically rolls back and a SQL_CODE of -752 is returned, which indicates that the connection could not be changed.

Stored procedures

A stored procedure is a compiled program stored in a database that can execute SQL statements. In Net.Data, stored procedures are called from Net.Data functions using a CALL statement. Stored procedure parameters are passed in from the Net.Data function parameter list. You can use stored procedures to improve performance and integrity by keeping compiled SQL statements with the database server. Net.Data supports the use of stored procedures with DB2 through the SQL language environment. For DB2 in particular, Net.Data supports stored procedures returning one or more result sets.

This section describes following topics:

- “Stored procedure syntax”
- “Calling a stored procedure” on page 86
- “Passing parameters” on page 87
- For DB2 only: “Processing result sets from DB2 stored procedures” on page 87

Stored procedure syntax

The syntax used for stored procedures includes the FUNCTION statement, the CALL statement, and optionally a REPORT block.

```
%FUNCTION function_name ([IN datatype arg1, INOUT datatype arg2,
    OUT resultsetname, ...]) {
    CALL stored_procedure
[%REPORT [(resultsetname)] { %}]
...
[%REPORT [(resultsetname)] { %}]
[%MESSAGE %]}

%}
```

Where:

function_name

Is the name of the Net.Data function that initiates the call of the stored procedure

stored_procedure

Is the name of the stored procedure

datatype

Is one of the database data types supported by Net.Data as shown in Table 6 on page 86. The data types specified in the parameter list must match the data types in the stored procedure. See your database documentation for more information about these data types.

For DB2 only: *tablename*

Is the name of the Net.Data table in which the result set is to be stored (used only when the result set is to be stored in a Net.Data table. If specified, this parameter name must match the associated parameter name for *resultsetname*.

For DB2 only: *resultsetname*

Is the name that associates a result set returned from a stored procedure with a REPORT block and a table name on the function parm list, or both. The *resultsetname* on a REPORT block must match a *tablename* on the function parameter list.

Table 6. Supported Stored Procedure Data Types

CHAR	FLOAT	TIME
DATE	GRAPHIC	TIMESTAMP
DECIMAL	INTEGER	VARCHAR
DOUBLE	REAL	VARGRAPHIC
DOUBLEPRECISION	SMALLINT	

Calling a stored procedure

1. Define a function that initiates a call to the stored procedure.

```
%FUNCTION (DTW_SQL) function_name()
```

2. (Optional) Specify any IN, INOUT, or OUT parameters for the stored procedure including the result set name of any result sets that are returned from the stored procedure.

```
%FUNCTION (DTW_SQL) function_name (IN datatype  
arg1, INOUT datatype arg2,  
OUT resultsetname...)
```

3. Use the CALL statement to identify the stored procedure name.

```
CALL stored_procedure
```

4. For DB2: If the DB2 stored procedure is going to generate one result set, optionally specify a REPORT block to define how Net.Data displays the result set.

```
%REPORT {  
...  
%}
```

Example:

```
%FUNCTION (DTW_SQL) mystoredproc (IN CHAR(30) arg1 OUT mytable) {  
    CALL myproc  
    %REPORT {  
        ...  
        %ROW { ... %}  
        ...  
    %}  
%}
```

5. If the stored procedure is going to generate more than one result set:

- Specify the result sets as OUT parameters in the FUNCTION statement. The result sets are saved as local tables.

```
%FUNCTION (DTW_SQL) function_name (OUT tablename, ...)
```

- Optionally specify one or more REPORT blocks to define how Net.Data displays the result sets.

```
%REPORT[(resultsetname1)] {  
...  
%}
```

Example:

```
%FUNCTION (DTW_SQL) mystoredproc (IN CHAR(30) arg1, OUT table1, table2) {  
    CALL myproc  
    %REPORT (table1) {  
        ...  
        %ROW { ... %}
```

```

    ...
%}
%REPORT (table1) {
    ...
    %ROW { ... %}
    ...
%}
%}

```

Passing parameters

You can pass parameters to a stored procedure and you can have the stored procedure update the parameter values so that the new value is passed back to the Net.Data macro. The number and type of the parameters on the function parameter list must match the number and type defined for the stored procedure. For example, if a parameter on the parameter list defined for the stored procedure is INOUT, then the corresponding parameter on the function parameter list must be INOUT. If a parameter on the list defined for the stored procedure is of type CHAR(30), then the corresponding parameter on the function parameter list must also be CHAR(30).

Example 1: Passing a parameter value to the stored procedure

```

%FUNCTION (DTW_SQL) mystoredproc (IN CHAR(30) valuein) {
    CALL myproc
...

```

Example 2: Returning a value from a stored procedure

```

%FUNCTION (DTW_SQL) mystoredproc (OUT VARCHAR(9) retvalue) {
    CALL myproc
...

```

Processing result sets from DB2 stored procedures

You can return one or more result sets from a stored procedure. The result sets can be stored in Net.Data tables for further processing within your macro or processed using a REPORT block. If a stored procedure generates multiple result sets, you must associate a name with each result set generated by the stored procedure. This is done by specifying parameters on the FUNCTION statement. The name you specify for a result set can then be associated with a REPORT block or a Net.Data table, enabling you to determine how each result set is processed by Net.Data. You can:

- Have the result processed in Net.Data's default report style by not defining a report block for the result set.
- Associate a result set with a REPORT block to apply your own report style. In the REPORT block, you can use Net.Data variables, text processing statements like HTML or JavaScript, or other functions to specify how the report data is displayed in the browser.

Result sets are always stored in local tables so that another function in the macro can also access the data. For example, you can pass a Net.Data table to another function so that it can use the data for calculations and display the results based on those calculations.

See "Guidelines and restrictions for multiple REPORT blocks" on page 64 for guidelines and restrictions when using multiple report blocks.

To return a single result set and use default reporting:

Use the following syntax:

```

%FUNCTION (DTW_SQL) function_name (OUT tablename) {
    CALL stored_procedure
%}

```

For example:

```
%FUNCTION (DTW_SQL) mystoredproc(OUT mytable1) {
    CALL myproc
%}
```

To return a single result set and specify a REPORT block:

Use the following syntax:

```
%FUNCTION (DTW_SQL) function_name (OUT tablename) {
    CALL stored_procedure [(resultsetname)]
    %REPORT [(resultsetname)] {
        ...
    %}
%}
```

Example 1:

```
%FUNCTION (DTW_SQL) mystoredproc (OUT mytable1) {
    CALL myproc
    %REPORT {
        ...
        %ROW { ... %}
        ...
    %}
%}
```

Example 2:

```
%FUNCTION (DTW_SQL) mystoredproc () {
    CALL myproc
    %REPORT {
        ...
        %ROW { ... %}
        ...
    %}
%}
```

To return multiple result sets and display them using default report formatting:

Use the following syntax:

```
%FUNCTION (DTW_SQL) function_name (OUT tablename1, tablename2) {
    CALL stored_procedure
%}
```

Where no report block is specified.

For example:

```
%DEFINE DTW_DEFAULT_REPORT = "YES"
%FUNCTION (DTW_SQL) mystoredproc (OUT mytable1, mytable2) {
    CALL myproc
%}
```

To return multiple result sets and specify REPORT blocks for display processing:

Each result set is associated with its one or more REPORT blocks. Use the following syntax:

```
%FUNCTION (DTW_SQL) function_name (OUT tablename1, tablename2, ...) {
    CALL stored_procedure
    %REPORT (resultsetname1)
        ...
        %ROW { ... %}
        ...
    %}
    %REPORT (resultsetname2)
        ...
    %}
```

```

    %ROW { ... %}
    ...
%}

...
%}

```

For example:

```

%FUNCTION(DTW_SQL) mystoredproc (OUT mytable1, mytable2) {
    CALL myproc

    %REPORT(mytable1) {
        ...
        %ROW { ... %}
        ...
    %}

    %REPORT(mytable2) {
        ...
        %ROW { ... %}
        ...
    %}
%}

```

Example 2:

```

%FUNCTION(DTW_SQL) mystoredproc(OUT mytable4, OUT mytable3) {
    CALL myproc (mytable1, mytable2, mytable3, mytable4)
    %REPORT(mytable2) {
        ...
        %ROW { ... %}
        ...
    %}
    %REPORT(mytable1) {
        ...
        %ROW { ... %}
        ...
    %}
    %REPORT(mytable4) {
        ...
        %ROW { ... %}
        ...
    %}
%}

```

The result sets *mytable2*, *mytable1*, and *mytable4* are processed by their corresponding REPORT blocks, in that order, and are displayed as specified. The result sets *mytable4* and *mytable3* are stored into table variables for further processing. The result set *mytable3* will also be displayed using Net.Data's default report format after the three REPORT blocks are done processing.

SQL language environment example

The following example shows a macro with a DTW_SQL function definition that calls an SQL stored procedure. It has three parameters of different data types. The DTW_SQL language environment converts the character string values in each parameter to the correct internal format and passes each parameter by reference to the SQL stored procedure. When the SQL stored procedure completes processing, the updated internal representation is converted to a character string and placed in the corresponding parameter.

```

%{*****DEFINE BLOCK*****%}
DEFINE {
    MACRO_NAME      = "TEST ALL TYPES"
    DTW_HTML_TABLE = "YES"
    Procedure       = "NDLIB.TESTTYPE"
    parm1           = "1"                %{SMALLINT      %}
}

```

```

parm2      = "11"           %{INT           %}
parm3      = "1.1"         %{DECIMAL (2,1) %}
%}

%FUNCTION(DTW_SQL) CRTPROC(){
CREATE PROCEDURE $(Procedure)
( INOUT SMALLINT,
  INOUT INT,
  INOUT DECIMAL(2,1))
EXTERNAL NAME $(Procedure) LANGUAGE C SIMPLE CALL
%MESSAGE{
  default : "$(DTW_DEFAULT_MESSAGE) : continuing.<br />": continue
%}
%}

%FUNCTION(DTW_SQL) myProc
( INOUT SMALLINT parm1,
  INOUT INT       parm2,
  INOUT DECIMAL(2,1) parm3){
CALL $(Procedure)
%}

%HTML(REPORT) {
<head>
<title>Net.Data : SQL Stored Procedure: Example '$(MACRO_NAME)'. <?TITLE>
</head>
<body bgcolor="#bbffff" text="#000000" link="#000000">
<p><p>
Calling the function to create the stored procedure.
<p><p>
@CRTPROC()
<hr>
<h2>
Values of the INOUT parameters
prior to calling the stored procedure:<p>
</h2>
<b>parm1 (SMALLINT)</b><br />
$(parm1)<p>
<b>parm2 (INT)</b><br />
$(parm2)<p>
<b>parm3 (DECIMAL)</b><br />
$(parm3)<p>
<p>
<hr>
<h2>
Calling the function that executes the stored procedure.
</h2>
<p><p>
@myProc(parm1,parm2,parm3)
<hr>
<h2>
Values of the INOUT parameters after
calling the stored procedure:<p>
</h2>
<b>parm1 (SMALLINT)</b><br />
$(parm1)<p>
<b>parm2 (INT)</b><br />
$(parm2)<p>
<b>parm3 (DECIMAL)</b><br />
$(parm3)<p>
</body>
%}

```

System language environment

The System language environment supports executing commands and calling external programs.

Issuing commands and calling programs

To issue a command, define a function that uses the System (DTW_SYSTEM) language environment that includes a path to the command to be issued in an EXEC statement. For example:

```
%FUNCTION(DTW_SYSTEM) sys1() {  
    %EXEC { /QSYS.LIB/ADDLIBLE.COM LIB(MYLIBRARY) %}  
%}
```

You can shorten the path to executable objects if you use the EXEC_PATH configuration variable to define paths to directories that contain the objects (such as, commands and programs). See “EXEC_PATH” on page 16 to learn how to define the EXEC_PATH configuration variable.

Example 1: Calls a program

```
%FUNCTION(DTW_SYSTEM) sys3() {  
    %EXEC { /QSYS.LIB/MYLIB.LIB/MYPGM.PGM %}  
%}
```

Tip: When calling programs, use the Direct Call language environment because it is more efficient and easier to use.

Passing parameters to programs

There are two ways to pass information to a program that is invoked by the System (DTW_SYSTEM) language environment, directly and indirectly.

Directly

Pass parameters directly on the call to the program. For example:

```
%DEFINE INPARAM1 = "SWITCH1"  
  
%FUNCTION(DTW_SYSTEM) sys1() {  
    %EXEC{  
        /QSYS.LIB/NETDATA.LIB/RPGCALL1.PGM ('$(INPARAM1)' 'literalstring')  
    %}  
%}
```

The Net.Data variable INPARAM1 is referenced and passed to the program. The parameters are passed to the program in the same way the parameters are passed to the program when the program is called from the command line. The parameters that are passed to the program using this method are considered input parameters, and any modification to the values are not reflected back to Net.Data (the parameters passed to the program can be used and manipulated by the program, but changes to the parameters are not reflected back to Net.Data).

Indirectly

Pass parameters indirectly, by using *environment variables*. Environment variables are character strings of the form "name=value" that are stored in an environment space outside of the program. The strings are stored in a temporary space associated with the process.

When Net.Data calls a DTW_SYSTEM language environment function, the language environment stores any function parameters that are input (IN) or input/output (INOUT) in the environment space prior to executing the statement within the %EXEC block. After the successful completion of the statement, the DTW_SYSTEM language environment determines whether there are any output (OUT or INOUT) function parameters. If so, the language environment retrieves the value corresponding to the function parameter from the environment space and updates the function parameter value with the new value. When Net.Data gets control, it in turn updates all OUT or INOUT parameters with the new values obtained from the DTW_SYSTEM language environment.

Set and retrieve environment variables using the APIs described in Table 7:

Table 7. Environment Variable APIs

ILE Programming Language	To retrieve, use...	To set, use...
C, C++	getenv()	putenv()
CL(1), RPG, COBOL	QtmhGetEnv()(2)	QtmhPutEnv()(2)

1. For OS/400 V3R7 and on, you can also use the CHGENVVAR and ADDENVVAR CL commands to set an environment variable.
2. QtmhGetEnv() and QtmhPutEnv() are shipped as part of the IBM HTTP Server.

You can pass Net.Data tables to a program called by the System language environment. The program accesses the values of a Net.Data macro table parameter by their Net.Data name. The column headings and field values are contained in variables identified with the table name and column number. For example, in the table myTable, the column headings are myTable_N_j, and the field values are myTable_V_i_j, where *i* is the row number and *j* is the column number. The number of rows and columns for the table are myTable_ROWS and myTable_COLS.

System language environment example

The following example shows a macro that uses the System language environment to issue the Send Break Message (SNDBRKMSG) command to all workstation message queues. The text of the message to be sent is constructed from form data (msgToSend).

```
%FUNCTION(DTW_SYSTEM) sndbrkmsg () {
  %EXEC { /QSYS.LIB/SNDBRKMSG.COMD MSG('$(msgToSend)') TOMSGQ(*ALLWS) %}
%}
%HTML(sndbrkmsg) {
  @sndbrkmsg()
%}
```

Chapter 7. Transaction management with persistent macros

Net.Data provides support for transaction processing with persistent macros. A *persistent macro* is a macro that contains built-in functions that enable the macro to run as part of a persistent CGI process in the Web server. This means that multiple blocks of a macro, or multiple macros, can run as part of a single logical transaction.

With non-persistent macros, Net.Data treats each macro invocation as one complete transaction. This means that after each response is sent to the browser, databases are committed, resources are released, and everything is set to an initial state. The next invocation of the same macro results in re-establishing the state of the application based on information passed into the macro as form data or information in the macro itself. There is no capability to save macro variables across invocations, to rollback database changes without explicitly undoing the changes made, or to treat database changes across multiple browser sessions as one complete transaction.

With persistent macros, as an application developer, you can build your application at a transaction level, invoking one or more macros while maintaining a persistent connection. This means variable data is persistent across invocations, so that you no longer need to pass information (such as user login ID) between macro invocations as hidden variables. This includes Net.Data table variables, which cannot be passed across invocations in non-persistent macros. Most important, the application can rollback all the work if the user decides to cancel out while in the middle of a transaction.

See “Invoking a persistent macro” on page 34 to learn about invoking persistent macros.

This chapter describes the following topics:

- “About persistent macros”
- “Defining a transaction” on page 94
- “Example of a persistent macro” on page 99

About persistent macros

When using persistent macros, Net.Data runs in a special persistent CGI process of the Web server, receives input through standard input and environment variables, and provides data through standard output. However, after the output is returned to the Web server, The Web server does not have to terminate the Net.Data process. Instead, the process remains active, waiting for a response from the user through the Web browser. Because the process does not terminate, Net.Data can maintain state information for the macro and can leave transactions open.

Net.Data communicates to the Web server that it wants to run in a persistent CGI process by sending the server a new HTTP header. Support for the new header, “Accept-HTSession”, has been added to the AS/400 HTTP Server in version 4, release 3 (V4R3). Net.Data decides what HTTP headers to send to the server when it sends its first output, because the headers must precede the output. This has the following implications to you when developing a persistent macro:

- Net.Data must know at the time the first output is generated from the macro whether or not this is to be a persistent macro.
- Using the new persistent macro built-in functions, you must specify the macro is persistent before any output is generated.

These restrictions will be noted in the documentation that follows.

The characteristics of persistent Net.Data processes are very similar to those of standard Net.Data processes with the following exceptions:

- They run in a pseudo-connection-oriented environment. The connection between Net.Data and the Web server is persistent, but the connection between the browser and the Web server still has no connection.
- They can have long running transactions. Because a single Net.Data process can span multiple browser requests, transactions can be left open and committed or rolled-back as appropriate based on subsequent browser requests or error conditions.
- A persistent Net.Data process can consume more system resources because it can remain active for potentially long periods of time. Care must be taken in the management of those resources.
- Portability is reduced because the Web server must contain support for persistence.

Defining a transaction

A transaction can span one HTML block, multiple HTML blocks, or multiple macros. When you specify that you want the macro to be persistent within a transaction, you need to define the start and end of the transaction, as well as which HTML blocks are included in the transaction. Net.Data provides built-in functions that help you complete the following persistent macro tasks:

- “Starting a transaction”
- “Specifying the macro HTML blocks in a transaction” on page 95
- “Ending a transaction” on page 97
- “Defining the scope of a variable in a transaction” on page 98
- “Specifying COMMIT and ROLLBACK in a transaction” on page 98

Starting a transaction

You start a transaction by indicating to Net.Data that a macro is persistent in your macro before any output is sent to the browser. Net.Data then sends a special HTTP header to the Web server to tell it that the macro requires persistent CGI support.

To start a transaction:

Use one of the following methods in the macro before any output is sent to the Web browser:

- Call the DTW_STATIC() built-in function.

The DTW_STATIC() function tells Net.Data that the current macro is persistent.

Syntax: @DTW_STATIC (["*timeout*"])

Where *timeout* is an optional parameter that specifies the number of seconds the Web server should wait for a response from the browser before ending the transaction.

Example:

```
@DTW_STATIC("60")
%DEFINE {
    var1 = "val1"
    var2 = "val2"
    %}
...

%HTML(input){
    ...
    %}

%HTML(report){
    ...
    %}
```

A timeout value of 60 seconds is specified for this transaction. If a response is not received within 60 seconds from the browser, the Web server ends the transaction. This does not affect the current page on the browser. However, the next page, which would have been part of the transaction, is now part of a new transaction.

- Define a variable with the STATIC attribute.

Syntax: %DEFINE(STATIC) var1 = "val1"

Example:

```
%DEFINE(STATIC) var1 = "val1"
%DEFINE var2 = "val2"
...
%HTML(input){
...
%}
%HTML(report){
...
%}
```

A statically defined variable keeps its value throughout a transaction, which can span multiple Net.Data invocations.

Specifying the macro HTML blocks in a transaction

You define which HTML blocks are a part of your transaction by using an identifier, called the *transaction handle*, in the URL request that invokes the HTML blocks. There are three steps in defining and using a transaction handle:

1. Define the transaction handle in your macro.
2. Call the DTW_ACCEPT built-in function to pass the handle name to Net.Data and the Web server.
3. Specify the handle in the URL request to invoke your next HTML block.

To define a transaction handle:

1. Define a variable for the transaction handle in the DEFINE section. For example:

```
%DEFINE handle=""
```

2. Optionally generate a unique transaction handle by specifying the DTW_RTVHANDLE() built-in function in the DEFINE section.

Syntax: @DTW_RTVHANDLE(*handle_name*)

Example:

```
@DTW_STATIC()

%DEFINE handle = ""
@DTW_RTVHANDLE(handle)
```

The transaction handle can be any valid character string. However, the DTW_RTVHANDLE() function provides a measure of security by generating a unique transaction handle, preventing others from invoking a macro which would run in your transaction.

To specify a transaction handle to Net.Data:

Specify the value of the transaction handle to Net.Data with the DTW_ACCEPT() built-in function. Because this handle is part of the information contained in the HTTP headers sent to the server, the DTW_ACCEPT() function must be called before any output is generated by the macro. Typically, it will be the first element in your HTML block.

Syntax: @DTW_ACCEPT(*handle_name*, [*timeout*])

Where *timeout* is an optional parameter that specifies the number of seconds the Web server should wait for a response from the browser before ending the transaction.

You can call DTW_ACCEPT() within an HTML block or outside of any HTML block. If the function is called outside of any HTML block, the transaction handle and the optional timeout values apply to all HTML blocks within the macro.

Example 1: Specifies a transaction handle for subsequent URL requests to run in this transaction

```
@DTW_STATIC()

%DEFINE handle = ""
@DTW_RTVHANDLE(handle)

%HTML(Block1){
@DTW_ACCEPT(handle)
...
%}
```

Important: When you call DTW_ACCEPT() as the first element in the HTML block, ensure that there is no white space between the line on which the %HTML statement is specified and the DTW_ACCEPT() call itself. Net.Data considers the white space as text to send to the browser, and issues an error because the DTW_ACCEPT() call is not found before data is sent to the browser.

Example 2: Specifies a transaction handle which applies to all HTML blocks in the macro

```
@DTW_STATIC()

%DEFINE handle = ""
@DTW_RTVHANDLE(handle)

@DTW_ACCEPT(handle)

%HTML(Block1){
...
%}

%HTML(Block2){
...
%}
```

To specify the handle when invoking an HTML block that is to run in your transaction:

After you have generated a transaction handle and called the DTW_ACCEPT() function, only URLs with that transaction handle can run in your transaction. The transaction handle must immediately follow the CGI program name in the URL.

Note: When entering the statements in your code, the URL should be on one line with no spaces, but it is split on two lines here for display purposes.

- HTML link:

```
<a href="http://server/Net.Data_invocation_path/transaction_handle/
filename/block/[?name=val&...]">any text</a>
```

- HTML form:

```
<form method=method
ACTION="http://server/Net.Data_invocation_path/transaction_handle/
filename/block/[?name=val&...]">any text</form>
```

- URL:

```
http://server/Net.Data_invocation_path/transaction_handle/
filename/block/[?name=val&...]
```

Parameters:

server Specifies the name of the Web server. If the server is the local server, you can omit the server name and use a relative URL.

Net.Data_invocation_path

The path and filename of the Net.Data executable file. For example, /cgi-bin/db2www/.

transaction_handle

Specifies which URLs are part of a transaction initiated by a Net.Data macro. The identifier is obtained by calling the DTW_RTVHANDLE built-in function and must follow the *Net.Data_invocation_path*.

filename

Specifies the name of the Net.Data macro. Net.Data searches for and tries to match this file name with the path statements defined in the MACRO_PATH initialization path variable. See "MACRO_PATH" on page 18 for more information.

block Specifies the name of the HTML block in the referenced Net.Data macro.

method Specifies the HTML method used with the form. METHOD=POST is recommended.

?name=val&...

Specifies one or more optional parameters passed to Net.Data.

Typically you will provide HTML links to these URLs or specify the URL on a form action tag in your macro.

Example 1: A typical HTML block with links to other macro invocations that run in the same transaction

```
@DTW_STATIC()
...
%define handle = ""
@DTW_RTVHANDLE(handle)

%html(report) {
@DTW_ACCEPT(handle)
...
<a href="/cgi-bin/db2www/${handle}/qsys.lib/mylib.lib/
  macros.file/pcgil.mbr/report2">continue</a><br />
<a href="/cgi-bin/db2www/${handle}/qsys.lib/mylib.lib/
  macros.file/pcgil.mbr/quit">quit</a><br />
%}
```

Example 2: A typical HTML block with a FORM ACTION link to another macro

```
@DTW_STATIC()
...
%define handle = ""
@DTW_RTVHANDLE(handle)

%html(input) {
@DTW_ACCEPT(handle)
...
<form method=post action="/cgi-bin/db2www/${handle}/qsys.lib/
mylib.lib/macros.file/pcgil.mbr/report2">
<p>What type of hardware do you want to see?
<menu>
<li><input type="radio" name="hardware" value="MON" checked />Monitors
<li><input type="radio" name="hardware" value="PNT" />Pointing devices
<li><input type="radio" name="hardware" value="PRT" />Printers
<li><input type="radio" name="hardware" value="SCN" />Scanners
</menu>
</form>
%}
```

Ending a transaction

You end a transaction by indicating to Net.Data that you no longer want your macro to be persistent.

To end the transaction:

Use the DTW_TERMINATE() built-in function to specify the end of a transaction. Like the DTW_ACCEPT() function, this function must be called before any output is generated by the macro and is typically specified as the first element in an HTML block. DTW_TERMINATE tells Net.Data that this invocation is the last invocation in the current transaction.

Syntax: @DTW_TERMINATE()

This function does not accept any parameters.

Example:

```
%html(quit) {  
@DTW_TERMINATE()  
...  
%}
```

Defining the scope of a variable in a transaction

You can decide what scope you want a variable to have in a transaction by specifying the scope as an attribute of the %DEFINE statement. You can specify

transaction scope

The variable scope is for the entire transaction.

single invocation scope

The variable scope is for a single Net.Data invocation.

To specify transaction scope for a variable:

Specify the attribute STATIC to indicate that the variable has transaction scope, meaning the value of the variable is saved across all invocations in a transaction. STATIC is the default for persistent macros. For example:

```
@dtw_static()  
%define(static) var1 = "val1"
```

To specify single invocation scope for a variable:

Specify the attribute TRANSIENT to indicate that the variable has single invocation scope, meaning the value of the variable will be re-initialized on each invocation. TRANSIENT is the default for non-persistent macros. For example:

```
@dtw_static()  
%define(transient) var1 = "val1"
```

In a persistent macro:

- All variables that *follow* the DTW_STATIC() call are STATIC if they are not explicitly defined as TRANSIENT.
- All variables that *precede* the DTW_STATIC() call are TRANSIENT if they are not explicitly defined as STATIC.

Specifying COMMIT and ROLLBACK in a transaction

In a non-persistent macro, a commit or rollback is done implicitly by Net.Data at the end of the macro invocation based on the success or failure of the invocation. With persistent macros, the commit or rollback is now done at the end of the transaction. However, because a transaction can span many macro invocations, you might want to commit or rollback changes incrementally within the transaction.

To commit pending changes during a transaction:

Specify the DTW_COMMIT() built-in function.

This function does not take any parameters and executes all changes pending in the transaction.

For example:

```
%html (report) {
@dtw_accept(handle)
...
%IF (action="Enter")
    @dtw_commit()
%ENDIF

%}
```

To rollback pending changes in the transaction:

Specify the DTW_ROLLBACK() built-in function.

This function does not take any parameters and backs out all changes pending in the transaction.

For example:

```
%html (report) {
@dtw_accept(handle)
...
%IF (action="Cancel")
    @dtw_rollback()
%ENDIF

%}
```

Example of a persistent macro

The following simple macro contains multiple HTML blocks that run in a single transaction:

```
@dtw_static()
%define a = "0"
%define(transient) b = "0"
%define handle = ""
@dtw_rtvhandle(handle)

%html(report) {
@dtw_accept(handle)
a = $(a)<br />
b = $(b)<br />
@dtw_add(a, "2", a)
@dtw_add(b, "2", b)
<a href="/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/report2">
click here to continue</a><br />
<a href="/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/quit">
click here to quit</a><br />
%}

%html(report2) {
@dtw_accept(handle)
a = $(a)<br />
b = $(b)<br />
@dtw_add(a, "2", a)
@dtw_add(b, "2", b)
<a href="/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/report3">
Click here to continue</a><br />
<a href="/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/quit">
Click here to quit</a><br />
%}
```

```

%html(report3) {
@dtw_accept(handle)
a = $(a)<br />
b = $(b)<br />
@dtw_add(a, "2", a)
@dtw_add(b, "2", b)
<a href="/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/quit">
Click here to quit</a><br />
%}

%html(quit) {
@dtw_terminate()
a = $(a)<br />
b = $(b)<br />
done
%}

```

Assuming the first call is to the HTML block report, Net.Data:

1. Calls the DTW_STATIC() function, which indicates that this macro is persistent.
2. Creates variable a as a STATIC variable because the default for persistent macros is STATIC.
3. Creates variable b as a TRANSIENT variable because it is explicitly defined with the TRANSIENT attribute.
4. Calls DTW_RTVHANDLE(), which generates a transaction handle and puts it in the variable handle.
5. Starts processing the HTML block report and calls DTW_ACCEPT(), which tells Net.Data what the transaction handle is for this transaction.
6. Finds output to send to the browser, which causes Net.Data to send the HTTP header to the Web server indicating a transaction is starting.
7. Displays the HTML page. The variables a and b both have a value of 0.

After the first page output is sent to the browser, users can choose to either continue with the transaction or quit. If they choose to continue, the Web server invokes URL:

```
/cgi-bin/db2www/$(handle)/WWW/MACROS/pcgi1.dtw/report2
```

The Web server recognizes the transaction handle as the one specified by Net.Data in the HTTP header. It invokes Net.Data as a persistent CGI program, which means the macro invocation is part of the current transaction.

When the HTML block report2 is invoked, Net.Data:

1. Calls the DTW_STATIC() function, which indicates this macro is persistent.
2. Recognizes that variable a is a STATIC variable and keeps the current value rather than re-initializing it to 0.
3. Recognizes that variable b is a TRANSIENT variable, creates a new instance of the variable, and initializes it to 0.
4. Calls DTW_RTVHANDLE(), which generates a transaction handle and puts it in the variable handle.
5. Starts processing the HTML block report2 and calls DTW_ACCEPT(), which tells Net.Data what the transaction handle is for this transaction.
6. Finds output to send to the browser, which causes Net.Data to send the HTTP header to the server indicating a transaction is continuing.
7. Displays the HTML page. Variable a will have a value of 2 and variable b will have a value of 0. The value of variable a is saved from the previous invocation because it is a static variable. The value of variable b is reset to 0.

After the second page is sent to the browser, the user can choose to either continue with the transaction or quit. If they choose to quit, the Web server invokes the following URL:

```
/cgi-bin/db2www/${handle}/WWW/MACROS/pcgi1.dtw/quit
```

The Web server recognizes the transaction handle as the one specified by Net.Data in the HTTP header, and invokes Net.Data as a persistent CGI program, which means the macro invocation is part of the current transaction.

When the HTML block `quit` is invoked, Net.Data:

1. Calls the `DTW_STATIC()` function, which indicates this macro is persistent.
2. Recognizes that variable `a` is a `STATIC` variable and keeps the current value rather than re-initializing it to 0.
3. Recognizes that variable `b` is a `TRANSIENT` variable, creates a new instance of the variable, and initializes it to 0.
4. Calls `DTW_RTVHANDLE()`, which generates a transaction handle and puts it in the variable `handle`.
5. Starts processing the HTML block `quit` and calls `DTW_TERMINATE()`, which tells Net.Data that this is the last invocation in this transaction.
6. Finds output to send to the browser, which causes Net.Data to send the HTTP header to the server indicating a transaction is ending.
7. Displays the HTML page. Variable `a` has a value of 4 and variable `b` has a value of 0.
8. Cleans up all variables and other resources that have transaction level scope, because the `DTW_TERMINATE()` call has been executed.

Chapter 8. Improving performance

Improving performance is an important part of tuning your system. This chapter discusses strategies for improving the performance of Net.Data. The following topics are discussed:

- “Net.Data caching of macros”
- “Optimizing the language environments”

In addition, ensure that your Web server has been properly tuned. The performance of your Web server has a direct effect on response time, independently of how fast Net.Data processes a macro or direct request.

Net.Data caching of macros

With Net.Data for OS/400, macro caching is enabled by default and is used to improve throughput and reduce CPU utilization. When macro caching is enabled, preprocessed macros are cached in memory when the macros are first invoked. These preprocessed versions are then available for reuse, thereby eliminating the costs associated with reading in the macros from HFS and processing them each time they are requested. The cached version of a macro is available to a requestor that has read permission for the file containing the macro.

The amount of memory that the preprocessed version of the macro uses is approximately twice the size of the macro itself. You can control the amount of memory that will be used for the caching of macros by using the caching configuration variable. For more information on using this variable, see “DTW_MACRO_CACHE_SIZE” on page 8.

Optimizing the language environments

The following sections describes techniques you can use to improve performance when using the Net.Data-provided language environments.

- “REXX language environment”
- “SQL language environment” on page 104
- “System language environment” on page 104

REXX language environment

Use the following tips to improve the performance of your Net.Data application:

- Combine your REXX programs where possible. Having fewer, larger programs provides better performance than more smaller programs because the REXX interpreter is initialized each time a REXX language environment function is called in the macro.
- Store the REXX program in an external file instead of including the REXX program inline in the Net.Data macro.
- For external REXX programs, reference the global variables on the command line in the %EXEC statement.
- Pass input-only parameters directly to a REXX program by defining global Net.Data variables and referencing the variables. For inline REXX programs, reference the global variables directly in your REXX source.
- To avoid the overhead of launching the REXX interpreter, consider using MACRO_FUNCTION blocks as an alternative to REXX programs.

SQL language environment

To learn about DB2 performance considerations, see *DB2 for IBM i SQL Programming*. This publication has a wealth of information, such as effectively using SQL indexes, improving performance of join queries, and improving performance when selecting data from more than two tables.

Use the following SQL language environment techniques to improve performance.

- Reduce the number of user IDs connecting to a database to avoid reconnecting to the database. The SQL language environment associates a user-profile and password to any remote connections to databases that it establishes. If the LOGIN and PASSWORD variables do not match the user-profile and password associated with an opened connection, the connection is closed and re-established, and the LOGIN and PASSWORD values are associated with the re-opened connection.
- Use the START_ROW_NUM and RPT_MAX_ROWS Net.Data variables to reduce the size of returned tables. If, on a SELECT SQL statement, the result set contains hundreds of records, return a subset of the result set back to the browser by using the START_ROW_NUM like a scrollable cursor and RPT_MAX_ROWS to limit the number of records returned. You should be aware that Net.Data reissues the query every time since there is no notion of state. However, you can use Net.Data support for persistent macros to store the result set in a Net.Data table that persists for the life of the transaction. See Chapter 7, “Transaction management with persistent macros,” on page 93 to learn more about persistent Net.Data macros.
- Consider using a stored procedure to handle complex database tasks. Using embedded SQL and understanding the structure of result sets reduces the overhead Net.Data uses to dynamically describe results. For more information on the performance trade-offs when using stored procedures, see the *DB2 Administration Guide*.

Note: Starting in OS/400 V4R2, the SQL engine has a prepared statement cache. Using the cache, the SQL engine stores away information about prepared statements, and keeps this information in system-wide storage. Then, when the same statement is executed again, even if its by a different user and a different job, the statement will run much faster. The system-wide prepared statement cache is part of normal SQL processing and requires no user action to configure or enable it. The cache may reduce any performance benefits that the static SQL might have over dynamic SQL.

System language environment

Pass input-only parameters directly to the program that the System language environment is invoking by defining global Net.Data variables and referencing the variables.

Net.Data for OS/400 has introduced a new language environment called Direct Call, which provides easier and more efficient interface for calling programs. Use the System language environment to issue commands; use the Direct Call language environment to call programs

Chapter 9. Serviceability features

The following sections describe tracing and error reporting features for Net.Data.

- “Net.Data trace log”
- “Net.Data error log”

Net.Data trace log

Tracing allows you to monitor Net.Data as it processes a macro. As Net.Data executes a macro, it will write out trace information that includes the functions being called, the parameters being passed (both input and output), and the errors that may be encountered.

Tracing is mainly used for service calls, but it can also be used by customers who are debugging their applications.

To enable tracing, you need to set where the trace log is stored and what level of trace data Net.Data needs to capture. This is done by setting the configuration variables `DTW_TRACE_LOG_DIR` and `DTW_TRACE_LOG_LEVEL`, respectively. For more information, see “`DTW_TRACE_LOG_LEVEL`” on page 14 and “`DTW_TRACE_LOG_DIR`” on page 14.

Net.Data logs trace records to the file, `netdata.trace`. In order for Net.Data to successfully write trace records to the trace log file, the user IDs under which Net.Data executes must have:

- Write authority on the log directory specified in the `DTW_TRACE_LOG_DIR` configuration variable.
- Execute authority on all directories in the path, including the log directory.

By default, trace records from the processing of a macro running in one thread are interleaved with trace records from the processing of a macro running in another thread. The default works well if you have total control over what URL is to be invoked against your Web site (as is the case in a development environment). However, in a high traffic environment, having the trace records of the processing of other macros interleaved with trace records that you may be interested in makes it hard to follow the flow of the processing of a macro. You can indicate to Net.Data whether trace records should or should not be merged into one file by setting the value of the configuration variable, `DTW_TRACE_MERGE_RECORDS` to `NO`, in which case Net.Data will log trace records to a file named `netdata.trace.XXXX`, where 'XXXX' is the process/thread identifier. For more information, see “`DTW_TRACE_MERGE_RECORDS`” on page 15.

If tracing is enabled, you can write your own trace messages to the Net.Data trace. To do this, simply pass the message as a parameter to the built-in function `DTW_LOG_TRACEMSG()`. See the *IBM Net.Data Reference* book for more information on how to use the built-in function `DTW_LOG_TRACEMSG()`.

Net.Data error log

Error logging allows you to have Net.Data errors logged to a file. If you are in the process of creating or enhancing a Net.Data application, you may want to log all error messages to ensure that known or expected errors are caught by message blocks. You may also want to track down unexpected errors. In production phase, log all uncaught error messages. Uncaught error messages indicate that something is wrong with your application. Periodically, check the error log to see if your application is working as expected. If error logging is disabled, the only way that you can know if something is wrong with your application is from your users.

To enable error logging, you need to set where the error log is stored and what level of error messages Net.Data needs to capture. This is done by setting the configuration variables `DTW_ERROR_LOG_DIR`

and `DTW_ERROR_LOG_LEVEL`, respectively. For more information, see “`DTW_ERROR_LOG_LEVEL`” on page 7 and “`DTW_ERROR_LOG_DIR`” on page 7.

Net.Data logs error records to the file, `netdata.error.log`. In order for Net.Data to successfully write error records to the error log file, the user IDs under which Net.Data executes must have:

- Write authority on the log directory specified in the `DTW_ERROR_LOG_DIR` configuration variable.
- Execute authority on all directories in the path, including the log directory.

If error logging is enabled, you can write your own error records to the Net.Data error log. To do this, simply pass the message as a parameter to the built-in function `DTW_LOG_ERRORMSG()`. See the *IBM Net.Data Reference* book for more information on how to use the built-in function `DTW_LOG_ERRORMSG()`.

Appendix A. Net.Data technical library

The Net.Data technical library is available from the Net.Data Web site at <http://www.ibm.com/systems/i/software/netdata/docs>.

Document	Description
<i>Net.Data Administration and Programming Guide</i>	Contains conceptual and task information about installing, configuring, and invoking Net.Data. Also describes how to write Net.Data macros, use Net.Data performance techniques, use Net.Data language environments, manage connections, and use Net.Data logging and traces for trouble shooting and performance tuning.
<i>Net.Data Reference</i>	Describes the Net.Data macro language, variables, and built-in functions.
<i>Net.Data Language Environment Interface Reference</i>	Describes the Net.Data language environment interface.
<i>Net.Data Messages and Return Codes</i>	Lists Net.Data error messages and return codes.

Related documentation

The following documents might be useful when using Net.Data and related products:

- *DB2 for IBM i SQL Programming*
- *IBM i Distributed Database Programming*

Additionally, OS/400 documentation and redbooks, including books about DB2, are available at the following URL:

<http://publib.boulder.ibm.com/series/>

Appendix B. Net.Data sample macro

The sample macro application below displays a list of employees names from which the application user can obtain additional information about an individual employee by selecting the employee's name from the list. The macro uses the SQL language environment to query the EMPLOYEE table for both the employee names and the information about a specific employee.

The macro uses an include file, which contains the DEFINE block for the macro.

```
%{***** Sample Macro *****)
*   FileName = sqlsaml.dtw
*   Description:
*       This Net.Data macro queries...
*       - The EMPLOYEE table to create a selection list of
*         employees for display at a browser
*       - The EMPLOYEE table to obtain additional information
*         about an individual employee
*
*****%}
%{*****
*   Include for global DEFINES -
*****%}
%INCLUDE "sqlsaml.hti"
%}
%{*****
*   Function: queryDB           Language Environment: SQL
*   Description: Queries the table designated by myTable and
*   creates a selection list from the result. The value of the
*   myTable is specified in the include file sqlsaml.hti.
*****%}
%FUNCTION(DTW_SQL) queryDB() {
  SELECT FIRSTNME FROM EMPLOYEE
%MESSAGE {
  -204: {<p><b>ERROR -204: Table EMPLOYEE not found. </b> </p>
        %} : exit
  +default: "WARNING $(RETURN_CODE)" : continue
  -default: "Unexpected ERROR $(RETURN_CODE)" : exit
%}

%REPORT {
<select name="emp_name">
%ROW{
<option>$(V1)</option>
%}
</select>
%}
%}

%{*****
*   Function: fname           Language Environment: SQL
*   Description: Queries the table designated by myTable for
*   additional information about the employee
*   identified by the variable emp_name.
*****%}
%FUNCTION(DTW_SQL) fname(){
SELECT FIRSTNME, PHONENO, JOB FROM EMPLOYEE WHERE FIRSTNME='$(emp_name)'
%MESSAGE {
  -204: "Error -204: Table not found "
  -104: "Error -104: Syntax error"
  100: "Warning 100: No records" : continue
```

```

+default: "Warning ${RETURN_CODE}" : continue
-default: "Unexpected SQL error" : exit
%}
%}
%{*****
* HTML block: INPUT Title: Dynamic Query Selection *
*
* Description: Queries the EMPLOYEE table to create a selection list *
* of the employees for display at the browser *
*****%}
%HTML(INPUT) {
<html>
<head>
<title>Generate Employee Selection List</title>
</head>
<body>
<h3>$(exampleTitle)</h3>
<p>This example queries a table and uses the result to create
a selection list using a <em>%REPORT</em> block. </p>
<hr />
<form method="post" action="report">
@queryDB()
<input type="submit" value="Select Employee" />
</form>
<hr />
</body>
</html>
%}

%{*****
* HTML block: REPORT *
* Description: Queries the EMPLOYEE table to obtain additional information *
* about an individual employee *
*****%}
%HTML(REPORT) {
<html>
<head>
<title>Obtain Employee Information</title>
</head>
<body>
<h3>You selected employee name = $(emp_name)</h3>
<p>Here is the information for that employee:
<pre>
@fname()
</pre></p>
<hr /><a href="input">Return to previous page</a>
</body>
</html>
%}

%{ End of Net.Data macro 1 %}

```

The file below is the include file used by the sample Net.Data application above.

```

%{***** Include File *****
* FileName = sqlsamp1.hti *
* Description: *
* This include file provides global DEFINES for the sqlsamp1.dtw *
* Net.Data macro. *
*****%}
#define {
emp_name = ""
reposition = sign
exampleTitle = "Sample Macro"
%}

%{ End of include file %}

```

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing 2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX	Language Environment
AS/400	MVS/ESA
DB2	Net.Data
DB2 Universal Database	OS/2
DRDA	OS/390
DataJoiner	OS/400
IBM	OpenEdition
IMS	

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

absolute path

The full path name of an object. Absolute path names begins at the highest level, or "root" directory (which is identified by the forward slash (/) or back slash (\) character).

ANSI American National Standard for Information Systems

API Application Programming Interface

applet A Java program included in an HTML page. Applets work with Java-enabled browsers, such as Netscape Navigator, and are loaded when the HTML page is processed.

BLOB Binary large object.

CGI Common Gateway Interface.

CLOB Character large object.

commitment control

The establishment of a boundary within the process that Net.Data is running under where operations on resources are part of a unit of work.

Common Gateway Interface (CGI)

A standardized way for a Web server to pass control to an application program and receive data back.

cookie A packet of information sent by an HTTP server to a Web browser and then sent back by the browser each time it accesses that server. Cookies can contain any arbitrary information the server chooses and are used to maintain state between otherwise stateless HTTP transactions.
Free Online Dictionary of Computing

current working directory

The default directory of a process from which all relative path names are resolved.

database

A collection of tables, or a collection of table spaces and index spaces.

database management system (DBMS)

A software system that controls the

creation, organization, and modification of a database and access to the data stored within it.

DATALINK

A DB2 data type that enables logical references from the database to a file stored outside the database.

data type

An attribute of columns and literals.

DBCLOB

Double-byte character large object.

DBMS

Database management system.

firewall

A computer with software that guards an internal network from unauthorized external access.

flat file interface

A set of Net.Data built-in functions that let you read and write data from plain-text files.

HTTP HyperText Transfer Protocol

hypertext markup language

A tag language used to write Web documents.

hypertext transfer protocol

The communication protocol used between a Web server and browser.

Internet

An international public TCP/IP computer network.

Intranet

A TCP/IP network inside a company firewall.

Java

An operating system-independent object-oriented programming language especially useful for Internet applications.

language environment

A module that provides access from a Net.Data macro to an external data source such as DB2 or a programming language such as REXX.

LOB Large object.

middleware

Software that mediates between an application program and a network. It manages the interaction between a client application program and a server through the network.

null A special value that indicates the absence of information.

path A search route used to locate files.

path name

Tells the system how to locate an object. The path name is expressed as a sequence of directory names followed by the name of the object. Individual directories and the object name are separated by a forward slash (/) or back slash (\) character.

persistence

The state of keeping an assigned value for an entire transaction, where a transaction spans multiple Net.Data invocations. Only variables can be persistent. In addition, operations on resources affected by commitment control are kept active until an explicit commit or rollback is done, or when the transaction completes.

port A 16-bit number used to communicate between TCP/IP and a higher level protocol or application.

registry

A repository where strings can be stored and retrieved.

relative path name

A path name that does not begin at the highest level, or "root" directory. The system assumes that the path name begins at the process's current working directory.

secure endpoint URL

Endpoint beginning with https

SSL Secure Sockets Layer

TCPIP Transmission Control Protocol/Internet Protocol

uniform resource locator

An address that names a HTTP server and optionally a directory and file name, for example: `http://www.ibm.com/software/data/net.data/index.html`.

unit of work

A recoverable sequence of operations that

are treated as one atomic operation. All operations within the unit of work can be completed (committed) or undone (rolled back) as if the operations are a single operation. Only operations on resources that are affected by commitment control can be committed or rolled back.

URL Uniform resource locator.

Web server

A computer running HTTP server software, such as Internet Connection.

wire All the underlying components that are responsible for physically sending or receiving a message on the web

XML eXtensible Mark-up Language

Index

A

- access rights
 - for language environments 71
 - for Net.Data files 21
- accessing DB2 82
- authentication, security 25
- authorization
 - security 26
 - specifying access rights to Net.Data files 21

B

- blanks, variable for removing extra 9
- blocks, macro 39

C

- caching macros, cache size 8
- calling
 - functions 55
 - Java applications 76
 - language environments 70
 - programs, Direct Call 72
 - programs, System 91
 - REXX programs 78, 79
 - stored procedures 85, 86
- CGI-BIN library, copying Net.Data program object 3
- COMMIT 84
- common errors when passing parameters 74, 84
- conditional
 - logic, IF blocks 65
 - variables 46
- configuration variable statements
 - configuring in the initialization file 5
 - description 5
 - DTW_DEFAULT_ERROR_MESSAGE 6
 - DTW_DEFAULT_MACRO 7
 - DTW_ERROR_LOG_DIR 7
 - DTW_ERROR_LOG_LEVEL 7
 - DTW_INTERPRET_HTTP_INPUT 7
 - DTW_JAVA_VMOPTIONS 8
 - DTW_MACRO_CACHE_SIZE 8
 - DTW_PAD_PGM_PARMS 9
 - DTW_PROCESS_REPORT_ON_ERROR 9
 - DTW_REMOVE_WS 9
 - DTW_REPORT_BLOCK_VARIABLE_SCOPE 10
 - DTW_RESRICT_PATH_SEARCH 10
 - DTW_SHOWSQL 10
 - DTW_SMTP_CCSSID 11
 - DTW_SMTP_CHARSET 11
 - DTW_SMTP_SERVER 12
 - DTW_SQL_CONNECTION_IDLE_TIME 12
 - DTW_SQL_ISOLATION 13
 - DTW_SQL_NAMING_MODE 14
 - DTW_TRACE_LOG_DIR 14
 - DTW_TRACE_LOG_LEVEL 14

- configuration variable statements
 - (continued)
 - DTW_TRACE_MERGE_RECORDS 15
- configuring Net.Data
 - access rights to Net.Data files 21
 - initialization file
 - configuration variable statements 5
 - creating 4
 - description 4
 - ENVIRONMENT statements 19
 - path statements 15
 - updating 4
 - overview 3
 - setting up language environments 20
- copying Net.Data program object
 - to CGI-BIN library 3
 - to multiple libraries 3
- creating initialization file 5

D

- data types
 - for Direct Call 72
 - for stored procedures 86
- database connection idle time 12
- declaration part, macro structure 37
- default reports
 - printing 61
 - specifying for stored procedures 87, 88
- DEFINE block
 - defining variables 43
 - description 39
- defining variables
 - DEFINE statement or block 43
 - HTML form SELECT, INPUT, and TEXTAREA tags 44
 - query string data 44
- Direct Call language environment
 - calling programs 72
 - common errors when passing parameters 74
 - overview 72
 - passing parameters 72
 - returning values from programs 74
 - supported data types 72
- DTW_ATTACHMENT_PATH 16
- DTW_DEFAULT_ERROR_MESSAGE 6
- DTW_DEFAULT_MACRO 7
- DTW_DEFAULT_REPORT 62
- DTW_DIRECTCALL 72
- DTW_ERROR_LOG_DIR 7
- DTW_ERROR_LOG_LEVEL 7
- DTW_INTERPRET_HTTP_INPUT 7
- DTW_JAVA_CLASSPATH 16
- DTW_JAVA_VMOPTIONS 8
- DTW_JAVAPPS 76
- DTW_MACRO_CACHE_SIZE 8
- DTW_PAD_PGM_PARMS 9

- DTW_PROCESS_REPORT_ON_ERROR 9
- DTW_REMOVE_WS 9
- DTW_REPORT_BLOCK_VARIABLE_SCOPE 10
- DTW_RESRICT_PATH_SEARCH 10
- DTW_REXX 78
- DTW_SHOWSQL 10
- DTW_SMTP_CCSSID 11
- DTW_SMTP_CHARSET 11
- DTW_SMTP_SERVER 12
- DTW_SQL 82
- DTW_SQL_CONNECTION_IDLE_TIME 12
- DTW_SQL_ISOLATION 13
- DTW_SQL_NAMING_MODE 14
- DTW_SYSTEM 91
- DTW_TRACE_LOG_DIR 14
- DTW_TRACE_LOG_LEVEL 14
- DTW_TRACE_MERGE_RECORDS 15
- DTW_UPLOAD_DIR 15, 34
- dynamically generating variable names 44

E

- encryption, network 25
- ENVIRONMENT statements
 - configuring in the initialization file 19
 - description 19
 - example 20
 - for user-defined language environments 5
 - language environment type 19
 - parameter list 19
 - service program 19
 - syntax 19
- environment variables 46
- error conditions, language environments 70
- executable variables 46
- executing commands 91
- executing SQL statements 82

F

- FFI_PATH 17
- files
 - specifying access rights to Net.Data 21
 - uploading 15, 34
- firewalls 23
- flat file functions 58
- footer information, REPORT block 61
- formatting data output 60
- forms
 - in Web pages to invoke Net.Data 33
 - invoking Net.Data 31, 35, 96
 - using the FILE input type 34
- FUNCTION block
 - calling functions 55
 - description 39

FUNCTION block *(continued)*

- formatting output 60
- identifier scope 42
- function calls
 - built-in 55
 - syntax 55
- functions
 - calling 55
 - calling stored procedures 85
 - defining 52
 - description 51
 - flat file 58
 - FUNCTION block syntax 52
 - general purpose 56
 - java applet 58
 - MACRO_FUNCTION block
 - syntax 52
 - math 57
 - persistent 59
 - string 57
 - table 58
 - user-defined 52
 - Web Registry 59
 - word 57

G

- general purpose functions 56
- global identifier scope 42

H

- header information, REPORT block 61
- hidden variables
 - conceal variable names 47
 - protecting assets 26
- HTML
 - blocks
 - description 40
 - example 59
 - invoking Net.Data 59
 - processing 60
 - FORM Submit button 60
 - forms
 - about 33
 - invoking Net.Data 31, 35, 96
 - SELECT, INPUT, and TEXTAREA tags, defining variables 44
 - generating in a macro 59
 - links
 - about 32
 - invoking Net.Data 31, 34, 96
 - tags for tables 61
 - unrecognized data as 60

I

- identifier scope 42
- IF blocks 65
- improving performance 103
- INCLUDE_PATH 17
- initialization file
 - configuration variable statements 5
 - creating 4, 5
 - description 4
 - ENVIRONMENT statements 19

initialization file *(continued)*

- format 4
- path statements 15
- updating 4
- invoking Net.Data
 - forms 31, 35, 96
 - HTML blocks 59
 - links 31, 34, 96
 - overview 31
 - URLs 35, 96
 - using CGI 31
 - with a macro 31

J

- java applet functions 58
- Java Application language environment
 - calling programs 76
 - overview 76
 - passing parameters 76
 - setting up 20

L

- language environments
 - calling 70
 - configuring ENVIRONMENT statements 19
 - configuring in the initialization file 19
 - Direct Call 72
 - examples 19
 - handling error conditions 70
 - Java Application 76
 - REXX 78
 - security 71
 - setting up 20
 - SQL 82
 - supported 70
 - System 91
 - variables 51
- links
 - in Web pages to invoke Net.Data 32
 - invoking Net.Data 31, 34, 96
- list variables 48
- looping, WHILE blocks 67

M

- macro request
 - examples 31
 - syntax 31
- MACRO_FUNCTION block
 - calling functions 55
 - syntax 52
- MACRO_PATH 18
- macros
 - anatomy 38
 - blocks 39
 - conditional logic 65
 - declaration part 37
 - DEFINE block 39
 - description 1
 - developing 37
 - FUNCTION block 39
 - functions 51

macros *(continued)*

- generating HTML 59
- HTML block 40
- identifier scope 42
- IF blocks 65
- looping 67
- navigation within and between 41
- persistent 93
- presentation part 37
- sample 38
- variables 42
- WHILE blocks 67
- math functions 57
- MESSAGE block
 - description 53
 - example 54
 - processing 53
 - scope 54
 - syntax 54
- miscellaneous variables 49
- multiple report blocks 62

N

- navigation, within and between
 - macros 41
- Net.Data
 - configuring 3
 - files, access rights 21
 - invoking 31
 - macros, developing 37
 - overview 1
 - security mechanisms 26
- Net.Data macros. See macros. 1
- Net.Data Program Object
 - copying to CGI-BIN libraries 3
 - copying to multiple libraries 3

P

- parts of a macro
 - declaration 37
 - presentation 37
- passing parameters
 - Direct Call language environment 72
 - Java Application language environment 76
 - REXX programs 79
 - stored procedures 87
 - System language environment 91
- path statements
 - configuring in the initialization file 15
 - DTW_ATTACHMENT_PATH 16
 - DTW_JAVA_CLASSPATH 16
 - DTW_UPLOAD_DIR 15
 - EXEC_PATH 16
 - FFI_PATH 17
 - INCLUDE_PATH 17
 - MACRO_PATH 18
 - protecting assets 26
 - update guidelines 15
- performance
 - optimizing language environments 103
 - REXX language environment 103

- performance *(continued)*
 - SQL language environment 104
 - System language environment 104
- persistent functions 59
- persistent macros 93
- printing, disabling for default reports 61
- processing result sets, stored
 - procedures 87
- protecting assets 23

R

- referencing variables 44
- REPORT block
 - stored procedures 88
- REPORT blocks
 - default reports 62
 - description 60
 - examples 62
 - formatting data output 60
 - guidelines for multiple 64
 - header and footer information 61
 - multiple 62
 - restrictions 64
 - scope 43
 - stored procedures 88
- report formats, customizing 61
- report variables 50
- reports
 - default 62
 - generating multiple with one function call 62
- result sets
 - multiple
 - default reports 88
 - guidelines and restrictions 64
 - processing, stored procedures 87
 - single 87
- RETURN_CODE variable 53, 70
- returning values from programs 74
- REXX language environment
 - calling programs 79
 - overview 78
 - passing parameters 79
- ROW block, identifier scope 43

S

- sample macro 109
- scope, identifier
 - FUNCTION block 42
 - global 42
 - macro 42
 - REPORT block 43
 - ROW block 43
- security
 - authentication 25
 - authorization 26
 - firewall 23
 - language environments 71
 - Net.Data mechanisms 26
 - network encryption 25
 - overview 23
 - specifying access rights 21, 71
- SQL
 - isolation configuration variable 13

- SQL *(continued)*
 - naming mode configuration
 - variable 14
- SQL language environment
 - common errors when passing parameters 84
 - executing SQL statements 82
 - overview 82
 - setting up 21
- SQL statements, executing 82
- SQLCODEs 70, 71
- starting Net.Data 31
- stored procedures
 - calling from macro 85
 - default reports 87, 88
 - multiple result sets 88
 - passing parameters 87
 - processing result sets 87
 - REPORT blocks 88
 - single result sets 87
 - steps 86
 - valid data types 86
- string functions 57
- System language environment
 - calling programs 91
 - issuing commands 91
 - overview 91
 - passing parameters 91

T

- table functions 58
- table processing variables 50
- table variables 49
- token sizes 42
- transaction processing 93
- TRANSACTION_SCOPE 84
- types, variable 45

U

- uploading files 15, 34
- URLs
 - defining variables 44
 - invoking Net.Data 35, 96
- user-defined functions 52
- user-defined language environments, ENVIRONMENT statements 5

V

- variables
 - conditional 46
 - configuration, statements
 - default macro enablement (DTW_DEFAULT_MACRO) 7
 - description 5
 - DTW_ERROR_LOG_DIR 7
 - DTW_ERROR_LOG_LEVEL 7
 - DTW_INTERPRET_HTTP_INPUT 7
 - DTW_JAVA_VMOPTIONS 8
 - DTW_PROCESS_REPORT_ON_ERROR 9
 - DTW_RESRICT_PATH_SEARCH 10
 - DTW_TRACE_LOG_DIR 14
 - DTW_TRACE_LOG_LEVEL 14

- variables *(continued)*
 - configuration, statements *(continued)*
 - DTW_TRACE_MERGE_RECORDS 15
 - e-mail SMTP CCSID (DTW_SMTP_CCSSID) 11
 - e-mail SMTP character set (DTW_SMTP_CHARSET) 11
 - e-mail SMTP server (DTW_SMTP_SERVER) 12
 - initialization file 5
 - macro cache size (DTW_MACRO_CACHE_SIZE) 8
 - padding parameters with blanks (DTW_PAD_PGM_PARMS) 9
 - removing extra blanks (DTW_REMOVE_WS) 9
 - SHOWSQL enablement (DTW_SHOWSQL) 10
 - SMTP character sets (DTW_SMTP_CHARSET) 11
 - SMTP server (DTW_SMTP_SERVER) 12
 - SQL connection idle time (DTW_SQL_CONNECTION_IDLE_TIME) 12
 - SQL isolation (DTW_SQL_ISOLATION) 13
 - SQL naming mode (DTW_SQL_NAMING_MODE) 14
 - variable scope variable (DTW_REPORT_BLOCK_VARIABLE_SCOPE) 10
- defining 43
- description 42
- dynamically-generated references 44
- environment 46
- executable 46
- generating names dynamically 44
- hidden 47
- language environment 51
- list 48
- miscellaneous 49
- referencing 44
- report 50
- scope 42
- table 49
- table processing 50
- token sizes 42
- types 42, 45

W

- Web Registry functions 59
- WHILE blocks 67
- white space, variable for removing extra 9
- word functions 57



Printed in USA