



Release notes

Welcome	1
Read the DWE Installation Guide	1
Getting started with Data Warehouse Edition	1
The DWE Tutorial	1
The DWE Design Studio	1
The Design Studio Help System	2
What's New in DWE 9.1.2	2
Migrating Design Studio workspace projects from DWE 9.1.0 to DWE 9.1.2	3
Known issues	4
Globalization	4
DWE Integrated Installation Program	4
DWE Configuration Tool	6
DWE Design Studio	7
DWE Administration Console	11
DWE Intelligent Miner	15
Documentation Updates	15
DWE Design Studio	15

Welcome

IBM® DB2® Data Warehouse Edition, Version 9.1.2 (DWE) combines the strength of DB2 Universal Database™ (DB2) with the powerful business intelligence infrastructure from IBM. Data Warehouse Edition provides a comprehensive business intelligence platform with tools that your enterprise and partners need to deploy and build next-generation analytic solutions.

Read the DWE Installation Guide

Before you start the installation of DB2 Data Warehouse Edition, Version 9.1, you must read the *DWE Installation Guide*. This book contains important installation instructions that are not available anywhere else. The book is available in PDF format from the DWE Quick Start and Documentation CD. It is also contained in the DB2 and DWE Information Center:

publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp

Getting started with Data Warehouse Edition

The DWE Tutorial

DWE provides a tutorial that shows you how you can use the components of DWE to create an end-to-end business intelligence solution. This tutorial is an excellent starting point for learning how to use DWE. It introduces you to the highlights of DWE so that you can get up and running quickly and easily.

Using the DWE tutorial, you can learn how to build and deploy a BI solution that extends the capabilities of a DB2 data warehouse. The tutorial is based on a new sample database that describes the operations of a fictional retail company, JK Superstore.

The DWE Design Studio

The DWE Design Studio provides an infrastructure and a set of integrated tools for developing Business Intelligence (BI) solutions. You can use these tools to build, populate, and maintain tables and other

structures for data mining and OLAP analysis in the context of a DB2 data warehouse. The Design Studio is built on the Eclipse workbench, which is a powerful development environment that you can customize.

The Design Studio includes the following tools and features:

- Integrated physical data modeling, based on Rational® Data Architect
- SQL Warehousing Tool for data flow and control flow design
- Data mining, exploration, and visualization tools
- Tools for designing OLAP metadata, MQTs, and cube models
- Integration points with WebSphere® DataStage® ETL systems

By integrating these tools, the Design Studio offers time to value and managed cost for warehouse-based analytics. All of the tools you need to support your BI applications are in one place.

The Design Studio Help System

The Design Studio includes a local help system, a set of sample files with a related tutorial, and an Eclipse “cheat sheet” that you can work through to get started.

- To use the help system, select **Help Contents** from the **Help** menu. Use the navigation tree on the left side of the help browser or the search facility above the tree to find the information that you need. The help system contains information about the SQL Warehousing Tool, OLAP, and data mining components.
- To open the sample projects, select **File > New > Example... > Data Warehousing Examples**. Select one of the samples and click **Next > Finish**. You can use these sample projects to work through the DWE Tutorial at your own pace.
- To run the cheat sheet, select **Cheat Sheets** from the **Help** menu, then select **Data Warehousing > Data Flow Cheatsheet** and click **OK**.

The cheat sheet walks you through the steps of creating a new data design project, a physical data model, a data warehouse project, and a data flow.

You can use your keyboard to open the Eclipse help system. In the Windows operating system, press F1 to open the Eclipse help window, whereas in the Linux environment, press Shift+F1.

What's New in DWE 9.1.2

Data Warehouse Edition Version 9.1.2 adds the following enhancements.

- The client tools, including Design Studio, are now available on 32-bit Linux.

Important: To prevent a problem that might cause DWE to fail to install on Linux, ensure that you follow the instructions in "Preparing your Linux or UNIX environment for DWE" in the DWE Installation Guide.

- The DWE Tutorial has been enhanced to support the clients on Linux.
- You can now install the data warehouse server and application server components on HP-UX and on new levels of Linux and Solaris:
 - HP-UX on Itanium Integrity Series (64-bit)
 - SUSE Linux Enterprise Server 10 on x86 (32-bit), AMD64 (64-bit), and EM64T (64-bit)
 - Sun Solaris 10 (64 bit) on UltraSPARC

To find the latest updates to supported hardware and software, see this document:

www.ibm.com/support/docview.wss?rs=3058&uid=swg2255099

- Version 9.1.2 includes three new DB2 Data Warehouse Editions: Starter Edition, Intermediate Edition, and Advanced Edition. See the announcement for DB2 DWE 9.1.2 for descriptions of these editions.

Additionally, the Base Edition has been renamed to the Enterprise Base Edition and expanded to include the SQL Warehouse Tool plug-in for the DWE Design Studio, the DWE Administration Console, and WebSphere Application Server.

Migrating Design Studio workspace projects from DWE 9.1.0 to DWE 9.1.2

In DWE 9.1.2, you can open a Design Studio workspace that you created in version 9.1.1. No migration is necessary.

Do not try to open a Version 9.1.0 workspace directly in Version 9.1.2 of the Design Studio. To migrate a Version 9.1.0 workspace, you need to complete three main steps:

1. Export all of the projects from the Version 9.1 workspace, including referenced projects, before installing Version 9.1.2 of the Design Studio.
2. Install the Version 9.1.2 Design Studio and create a new workspace.
3. Import the exported projects into the new workspace.

You can export and import projects in several different ways. The following detailed instructions explain the recommended approach, which uses the **Export as File System** option. This example assumes that you have a Version 9.1.0 data warehouse project named V91DW, which contains control flows, data flows, and SQL statements. The V91DW project depends on a data design project named V91DATA.

Note: The first procedure, which is recommended, assumes that the Version 9.1.2 Design Studio is not yet installed. If you have already upgraded to the Version 9.1.2 Design Studio on the same computer where Version 9.1.0 was installed, follow the second procedure.

To migrate the V91DW and V91DATA projects (before installing Version 9.1.2):

1. Start the Version 9.1.0 Design Studio and select your existing workspace.
2. In the Data Project Explorer, right-click the V91DW project and select **Export**.
3. In the Export window, select **File system** and click **Next**.
4. Select the V91DW and V91DATA projects and enter an existing directory in the **To directory** field.
5. Close the Version 9.1.0 Design Studio.
6. Install the Version 9.1.2 Design Studio.
7. Start the Version 9.1.2 Design Studio and enter a new workspace location in the Workspace Launcher. For example: C:\DWEDesignStudio\workspace912
8. Close the Welcome view, which opens by default when you create a new workspace.
9. From the **File** menu, select **Import**.
10. In the Import window, select **General**, then **Existing Projects into Workspace**. Click **Next**.
11. In the **Select root directory** field, browse to the directory where you exported the projects in step 4. A list of projects is displayed.
12. Optional: Check the **Copy projects into workspace** check box.
13. Click **Finish**. The projects are imported.

Repeat this process for all of the projects in your Version 9.1.0 workspace.

To migrate the V91DW and V91DATA projects (after installing Version 9.1.2):

1. Start the Version 9.1.2 Design Studio and enter a new workspace location in the Workspace Launcher. For example: C:\DWEDesignStudio\workspace912
2. Close the Welcome view, which opens by default when you create a new workspace.
3. From the **File** menu, select **Import**.
4. In the Import window, select **General**, then **Existing Projects into Workspace**. Click **Next**.

5. In the **Select root directory** field, browse to the Version 9.1.0 workspace directory. A list of projects is displayed.
6. Optional: Check the **Copy projects into workspace** check box.
7. Click **Finish**. The projects are imported.

Repeat this process for all of the projects in your Version 9.1.0 workspace.

Known issues

While every attempt is made to ensure the highest quality for DB2 Data Warehouse Edition, the following known limitations and problems remain. See the Known Issues section of the *DWE Installation Guide* for additional limitations that are not documented here.

Globalization

Issues exposed in versions 9.1.0 and 9.1.1:

In general, database objects and data (tables, columns, rows, and so on) can use non-English locales. Data flows, mining flows, and control flows that contain non-English data run successfully in both the Design Studio and the Administration Console. In the Design Studio, you can use non-English names for objects such as data flows, control flows, mining flows, variables, and application profiles.

Note the following limitations in globalization support:

- You cannot log into the DWE Administration Console with a user ID that consists of double-byte character set (DBCS) characters. (23880)
- You cannot install DWE to a directory with a non-English path name. This is a known issue in DB2 Universal Database and WebSphere Application Server. (24113)
- Source, execution, and target databases must have the same character encodings; otherwise data flows might fail with data conversion errors such as error -302. To work around this problem, make sure that all of the databases that are used during execution have been defined with the same coded character set identifier (CCSID). (24819)
- On Windows, after installing the Simplified Chinese or Traditional Chinese version of DWE Alphablox, you might see some English words remaining in the Start menu structure. (144887)

DWE Integrated Installation Program

Issues exposed in version 9.1.2:

- **Important:** To prevent a problem that might cause DWE to fail to install on Linux, ensure that you follow the instructions in "Preparing your Linux or UNIX environment for DWE" in the DWE Installation Guide.
- If you install DB2 DWE 9.1.2 on HP-UX and work with Cube Views, you might see the following error:
aCC runtime: Use of "-mt" must be consistent during both compilation and linking

The error occurs if you have the PHSS_31855 aC++ Runtime (IA: A.05.61, PA A.03.61) patch installed and you invoke the **db2mdapiclient** utility in Cube Views. To fix the error, remove the patch from the computer.

- If you have Query Patroller installed with DB2 8 FP10 or later, and you install DWE 9.1.x, then you must rebind the Query Patroller package to its DB2 8 databases after installing DWE. This happens because the DWE installation program updates the DB2 instance, and all existing package bindings are lost. To re-bind the Query Patroller package to all existing databases, take the following steps:

1. Log in as a user with DBADM authority.
2. Change to the proper directory. For example:

On Linux and UNIX:

INSTPATH/bnd

On Windows:

`INSTPATH\bnd`

where `INSTPATH` is the DB2 instance path, such as: `/home/db2inst1/sqllib`

3. Run the following commands:

```
db2 connect to dbname
db2 bind @qpserver.lst blocking all grant public
```

where `dbname` is the name of the database.

(145735)

- In a partitioned environment, when you install DWE Cube Views or DWE Intelligent Miner on an additional partition, you might receive an exception in the log file. This can happen when the DWE installation program cannot link the Cube Views or Intelligent Miner libraries to the DB2 instance for DWE. To work around the problem, take the following steps:
 1. Create a DB2 instance for DWE if one does not exist.
 2. If you installed Intelligent Miner, run the **idminstfunc** utility. This utility is in directory where Intelligent Miner is installed. The utility creates the links for the mining libraries.
 3. If you installed Cube Views, run the DB2 command **db2iupdt** to create the links for Cube Views libraries.
 4. Copy the `db2md_config.xml` file to the instance directory.

(146109)

- When you install Query Patroller, you can install the server and client components at the same time, or you can install the server first and then later install the client. However, if you install the client alone, and then later try to install the server, an error occurs. (146020)
- When uninstalling DWE on a computer on which DB2 9 and DB2 8 are both installed, you might get an error when it the DWE uninstallation program tries to uninstall DB2 8. The solution is to uninstall DB2 8 using the native DB2 uninstallation program. (145810)
- On Linux, when you uninstall DWE, the clients are always uninstalled, in addition to the components you selected to be uninstalled. (145930, 146164)

Issues exposed in versions 9.1.0 and 9.1.1:

- On Red Hat Enterprise Linux 4 (RHEL 4), to avoid errors in the Administration Console, after you install DWE, install the RHEL 4 Legacy Software Development Package. (144975)
- You cannot install DWE to a directory with a non-English path name. This is a known issue in DB2 Universal Database and WebSphere Application Server. (24113)
- When you uninstall DB2 Alphablox, some WebSphere variables may remain configured. The variables `WS_EAR_AlphabloxPlatform` and `WS_EAR_ApplicationStudio` should be removed with the WebSphere Administration Console before you re-install DB2 Alphablox. Alternatively, you can uninstall and reinstall WebSphere to remove those entries. (23925)
- When you run the DWE configuration tool on the application server, use the same connection and user information that you supplied to the installation program when you installed DB2.
- If you install the combined DB2 and DWE Information Center, and then run the installation again, the installation program will allow you to select the installable image for the DWE documentation, even though those files are installed. (23717)
- When you install on AIX®, you might see a panel with this error message: "For input string: "2268987392"". This occurs if you do not have enough disk space in the `/tmp` or `/usr/opt` directories for installing DWE. Make sure your computer meets the installation hardware and software requirements listed in the DWE Installation Guide. (24066)
- When you copy the extracted installation files to a different location by using a utility or command that does not copy "hidden" files by default (such as `xcopy`), those files will be missing. Starting the

installation from the target location then causes the installation program to fail with a Java™ exception. Make sure that your copy command writes all of the files to the target location before you start the installation process. (24593)

- If an error occurs when you install DWE, and then after you recover and run the installation again, the install program might indicate that the DWE components that you selected are already installed, when they are not. This is caused by an Installshield limitation. You can continue to install the components despite the incorrect message. (94282)
- Under certain circumstances, the summary window of the installation program might not accurately reflect the exact DWE components you selected to install or uninstall. The components you have actually selected on the feature selection window are installed, rather than what is listed on the summary window. (94390)
- On Solaris, in the window on which you specify the location of the CD images to be installed, the **Update** button does not work. (143438)
- On UNIX, if you uninstall the DWE data warehouse server (including the DB2 Server) and then reinstall them, you might get an error. The error occurs when the installation program attempts to create a DB2 instance for the DB2 server, but it cannot do so because the uninstallation leaves behind some DB2 processes. To solve the problem, after you uninstall the DB2 server, either kill the leftover DB2 processes manually or restart the computer. Then you can install the DWE data warehouse server using same DB2 instance user name. (143604)
- When you uninstall, make sure the passwords for the DB2 IDs and user accounts are current and valid. If any of them are expired or locked, the uninstallation will hang indefinitely. If this happens, stop the uninstallation process, change the passwords so they are valid, then restart the uninstallation. (94670)
- After uninstalling from Linux 64-bit or Solaris, some DB2 files might remain on the computer. Do not delete the files. Instead, you can remove the files using the DB2 rpm utility. (144870)
- When you migrate from DWE 9.1 to DWE 9.1.1, you might see a NullPointerException error in the dweinstall.log file even though the installation completely successfully. This is a known defect in ISMP and has no effect on DWE. (144815)

DWE Configuration Tool

Issues exposed in versions 9.1.0 and 9.1.1:

- If you install WebSphere on a 64-bit system, the DB2 instance in which you create the control and schedule databases for the DWE Administration Console must also be 64-bit. If you configure the Administration Console with a 32-bit DB2 instance while using WebSphere on a 64-bit system, you see an error message like this:

```
java.sql.SQLException: Failure in loading T2 native library db2jcct2
DSRA0010E: SQL State = null, Error Code = -99,999
```

The DWE installation program on Linux® creates a 32-bit instance in DB2. If you want to run WebSphere in a 64-bit system, you must create a 64-bit DB2 instance. Take these steps to start that process:

1. Stop DB2.
2. Log in as root.
3. Go to /opt/IBM/db2/V8.1/instance and run the command **./db2issetup** .
4. Start DB2 as the 64-bit instance.

When you run the DWE Configuration Tool, make sure you point to this 64-bit DB2 instance.(24335)

- When you specify local OS groups, only specify one group for each role. (23274) If you need more groups, you can add them using the WebSphere Application Server administration tool.
- After configuration, if you want to authorize additional users to the DWE Administration Console, you must also authorize them to WebSphere using the WebSphere Application Server administration tool. (144671)

- While configuring DWE, the SCHD0046E message: "The table SCHED_TREG in datasource jdbc/dwesched used by scheduler resource ETLScheduler (ETLScheduler) is missing or is not accessible" is displayed. You can safely ignore this message. (23758)
- On a Linux 64-bit computer, when you run the configuration tool, you might see this error message: "Checking required disk space requires file service native support". This message has no affect and you can continue with the configuration. (143398)
- In DB2, table space names are case sensitive. If you create a table space name but do not surround it with quotes, DB2 converts the name to all uppercase. Then, if you run the DWE configuration tool and specify that table space name in lowercase for the control database, you get an error. (143496)
- When you configure both WebSphere Application Server and DB2 Alphablox at the same time, the user ID that you specify when you set global security becomes the administrator for both WebSphere and Alphablox. (144074)
- When you configure the e-mail server for WebSphere Application Server, the configuration tool does not set up authentication. If you want authentication, go to the WebSphere Administration Console and select **Mail Providers > DWEMailProvider > Mail Sessions > DWEMail > Custom Properties**. Then add the property **mail.smtp.auth** with a value of **true**. (145158)
- On Windows, if you install and configure DWE on a computer that is on a domain, then the ID you use to run the configuration tool must be a domain user account with Administrator privileges. Otherwise the ID can be a local user.

When you set global security in the configuration tool, the user ID that you specify must be a local user with Administrator privileges; it cannot be a domain user account. It is not necessary to specify the user ID you are using to run the configuration tool. (144909)

DWE Design Studio

Issues exposed in version 9.1.2:

- In the table target operator, you cannot use variables for table and schema names when you perform an SQL Update or a Delete operation on a remote database. (145755)
- When the Design Studio and the database are running on the same machine and are using the same instance user, you cannot connect to a local database by using the the Type 2 JDBC driver. Workaround: Use the Type 4 JDBC driver by selecting **IBM DB2 Universal** from the drop-down menu in the connection wizard. (146369)
- If you install the DWE Base Edition (which provides the SQW and Cube View plug-ins) and later upgrade to the DWE Enterprise Edition (which also includes the mining plug-in) you cannot create mining flows for a Data Warehouse project that was created before the upgrade. Workarounds: (1) Create a new Data Warehouse project that corresponds to the one you created before the upgrade, and copy all of the files from that project to the corresponding new project. (2) Create a new Data Warehouse project, and reference it to the project that was created before the upgrade to the DWE Enterprise Edition. (146108)
- Be sure to set up the DB2 environment correctly before running the Design Studio on Linux so that DB2 commands, such as those that data flows and control flows invoke, can run successfully. If multiple instances of the Design Studio client exist, be sure that you are running the appropriate instance of the Design Studio client, and be sure that the environment has been set for that instance. You set the environment for an instance in the shell as follows:
 - `source ~db2inst2/sqllib/db2cshrc` (for `cs`h/`tc`sh)
 - `. ~db2inst2/sqllib/db2profile` (for `sh`, `ksh`, `bash`)
 (145762)
- After connecting to two different databases in the Database Explorer, importing the Metadata for one of them, and then comparing the two databases, the Metadata importing feature stops working. Work around: Before you try to import the Metadata from a previously compared database, disconnect from that database, and then reconnect to it. (146114, 146099)

- During deployment preparation for data warehouse applications on Linux platforms, you might see a number of SQL Parser debugging errors (System.out/System.err messages). You can deploy the application and run its processes successfully despite these errors. (146207)
- On Linux platforms, JDBC Type 2 driver connections to databases in the Database Explorer might fail after a successful initial connection. If you are unable to reconnect to a database, either create a new Type 4 connection for the database or restart the Design Studio and try the connection again. (145805)

Issues exposed in versions 9.1.0 and 9.1.1:

- An existing connection that is defined in the Database Explorer will fail if you change its JDBC driver setting from **Other** (a Type 2 Driver URL) to **IBM DB2 Universal** (a Type 4 Driver URL). The test connection succeeds but the actual connection fails with an exception error. To work around this problem, delete the existing connection and create a new connection with the correct driver type. (145027)
- When you create a connection to a Microsoft SQL Server database in the Database Explorer, you must append `SelectMethod=cursor` to the connection URL. For example:

```
jdbc:microsoft:sqlserver://9.99.99.99:1433;databaseName=DWESAMP;SelectMethod=cursor
```

The database connection will succeed without this change to the URL, but data flow execution will fail. (145162)

- After closing the Design Studio, a connection to the DB2 database might remain for the application `javaw.exe`; you can see this process in the Windows[®] Task Manager. When you restart the Design Studio, you will see a "workspace in use" error. Ending the `javaw.exe` task in the Task Manager frees the workspace so that you can restart the Design Studio successfully. This problem occurs when the Design Studio is closed while a metadata replication process is in progress. To avoid this problem, always wait until metadata replication has finished before closing the Design Studio. (18984)
- On Windows 2000 Server or Windows 2003 Server, when you view the Welcome page of the Design Studio in Microsoft[®] Internet Explorer, the viewlets will not run if the Security Setting level in Internet Explorer is set to High. The viewlets will run correctly if you set the level to Medium or lower. (24551)
- When you run a data flow from the Design Studio and the target table belongs to a remote database, the execution log does not contain detailed statistics messages that are returned by the target database. The run profile summary contains the expected information messages that state whether the data flow succeeded or failed, but the execution log is empty. (22623)
- Connecting to CVS from the Design Studio: In order to use a Windows CVS server to work with Design Studio projects, you must install CVSNT V2.0.41. The repository paths set in CVSNT must be identical to the repository paths set in Design Studio. (23976)
- Long path names can cause control flows to fail during execution because of operating system limits on the maximum length of a path. These limits apply to control flows that you run in both the DWE Design Studio and the DWE Administration Console. The length of a control flow path is influenced by elements such as the installation path and the activity name. Using long names for any of these components can cause the maximum path length to be exceeded and the control flow to fail. Use good judgment when selecting installation paths and names. (24207)
- SCM plugin compatibility: You might need to use the Resource perspective (or other perspectives) and the Navigator view for specialized source control management (SCM) operations; check the documentation that is provided by the SCM plugin vendor. In some cases, the BI and Data perspectives and their views are not completely compatible with SCM plugins. For example, the different CVS perspectives offer specialized support and views for CVS operations, and it is more convenient to use those perspectives. In other cases, a manual refresh of the projects in the Data Project Explorer might be necessary. (24246)
- SQL support in subflows: Operators that are defined inside subflows must generate SQL that conforms to the SQL-92 standard. For example, if a select list operator in a subflow has a reference to a SQL-99 OLAP function, the generated code will be incorrect and the data flow will fail to run. This problem does not apply to operators that are used directly in data flows. (24438)
- Limitations of the SQL replication source operator:

- The Replica and CCD replication types are not supported in this release. (21723)
- A SQL replication source operator must use a replication target database server that is installed on the same computer as the SQL Warehousing Tool runtime environment. In other words, the SQL execution database that is used for the data flow must belong to a DB2 instance that is running on the same computer as either the DWE Administration Console or the DWE Design Studio. If this configuration is not used, the data flow that contains the SQL replication source operator will fail to run. (24350)
- Visibility of execution plan graph (EPG) files: EPG files are not visible in the Data Project Explorer. To see these files inside the project tree, switch to the Navigator view. The EPG files are inside the generated-code folder. (23301)
- The Database Explorer does not automatically display newly cataloged DB2 databases without manual intervention: If the Design Studio is open and you catalog a database in the DB2 system on the client machine (either a local database or a remote database), you must perform one of the following steps to display the new database in the Database Explorer:
 - Restart the Design Studio.
 - Use the Database Explorer to connect to the local database.

Note: The DB2 database that you use to execute data flows and control flows must first be cataloged locally even when the database instance is in a remote machine.

To connect to the database, do these recommended steps:

1. Right-mouse click in the Database Explorer and click **New Connection**. The New Connection dialog is displayed.
 2. In the **Select a database manager** field, select the appropriate version of DB2.
 3. In the **JDBC driver** field, select **Other**.
 4. Type the connection parameters in the appropriate fields to complete the New Connection dialog. (24006)
- Delta DDL generation in the compare and sync editor: You might experience problems with DDL generation after comparing OLAP objects in a physical data model in the Data Project Explorer with a corresponding set of objects in a database in the Database Explorer. In some cases the script that is generated to deploy the changes to the database might not execute successfully. If you encounter this problem, use an alternative procedure to deploy the changed OLAP objects to the database:
 1. Drop any objects from the target database that are the same as objects in the physical data model.
 2. Export the changed OLAP objects from the physical data model to a file, using the **File > Export** option. Select the **OLAP Metadata** option in the first page of the Export wizard.
 3. Import the OLAP objects from the file into the target database, using the **File > Import** option. Select the **OLAP Metadata** option in the first page of the Import wizard. (00499, 00500, 00523)
 - When you are working with a physical data model in the Data Project Explorer and you try to replace a table with an alias on the same table, you cannot select the columns to create the attributes. To solve this problem, you can create the alias in the database and reverse engineer the alias. (00525)
 - The following problem applies only to DB2 Version 8 Fix Pack 11 and earlier versions. This problem is fixed in DB2 Version 8 Fix Pack 12 and all later versions.

Mining flow execution aborts with an "abnormally terminated" error message: In certain cases, the execution of a correctly defined mining flow can fail with an "abnormally terminated" error message issued by DB2. Before this message is displayed, execution of the flow typically hangs for a few minutes. This problem can happen if two or more Intelligent Miner™ user-defined functions (UDFs) are called during processing of one and the same SQL statement. For example, when you have three quality extractor operators and connect the three output ports with the input ports of the same table join operator, the resulting SQL statement calls three DM_getQualities() functions and might fail (function calls contained in SQL views that are referenced by an SQL statement also count). Combining

the outputs of several extractor operators through a union or table join operator is a typical scenario where you might see this problem, but it can also occur with other mining operators.

As a workaround, append a data station operator to the output port of each mining operator that produces such an Intelligent Miner UDF (for example, append the data station operator to those extractor operators that cause the problem). Select `TEMPORARY_TABLE` as the station type for the data station operator, but do not check the **Pass through** option. In this way, the data is buffered and each of the resulting SQL statements calls only one Intelligent Miner UDF.(24594)

- A regression model that uses classification as a function cannot be used or opened in the Design Studio. This kind of model is never created by Intelligent Miner. This problem only occurs with models that are imported from other data mining vendors.
- The **Deploy job metadata to server** operation from the DataStage Servers view fails if you try to import an XML file from a directory location that has spaces in its path. To work around this problem, move the XML file that you need to import to a location that does not have spaces. (24700)
- In the Analyze Model panel, the description text in the text area at the bottom of the panel is not displayed in the non-English versions of the product. An error message is displayed instead. (dwe00000587)
- In the SQL Expression Builder, when you double-click a ByteArray variable, the text inserted in the **SQL Text** field is in the following form:
`x'${GRP/ByteArray}'`

At runtime, when the variable is replaced with a value, the DB2 database treats this value as a binary value because of the leading x. The data type of the column that results from the SQL expression is incorrectly changed to BLOB with length 0. You must manually alter the data type of the column to CHAR FOR BIT DATA and specify the appropriate length. This problem does not occur when ByteArray variables are used in SQL conditions. (144608)

- Generating the code for a data flow or mining flow automatically saves it. Accordingly, if you are working in DWE Version 9.1.1 and you generate the code for a data flow or mining flow that was created in DWE Version 9.1, you can no longer use that flow in DWE Version 9.1. Be sure to make copies or a backup of all of the data flows and mining flows that you create in DWE Version 9.1 that you want to use in both DWE Version 9.1 and DWE Version 9.1.1 before you generate the code for them in DWE Version 9.1.1. (141928)
- When you create a new connection to a Microsoft SQL Server database in the Database Explorer, and then create a new physical data model using that connection, the tables of the SQL Server database in the new physical data model are missing the CHAR and VARCHAR data types. If you use these tables as the source or target tables of a data flow, the data flow fails during the code generation process because of the missing data types. To resolve this issue, use the Microsoft SQL Server 2005 JDBC Driver version 1.1 and the Microsoft SQL Server 2000 Driver for JDBC Service Pack 3 that is dated 4/21/2004. You can download these drivers from the Microsoft web site. (145163)
- When you specify the directory of the execute command to execute a data flow that has been deployed as a DataStage executable, you need to surround the path to the directory with quotation marks if the path contains one or more spaces. For example: "C:\Program Files\IBM\dwe\Client\sqwtools\eclipse\plugins\com.ibm.datatools.etl.dis_1.0.0\scripts". (145078)
- Do not use browser buttons to navigate pages in the DWE Design Studio. Use only the application buttons and links to navigate pages. If you use the browser back and forward buttons, you put the application in an undefined state and an error will occur. (94130)
- A Design Studio user who needs to reverse engineer a database that has been enabled for data mining must connect to the database with a DB2 user ID that has authority to read mining metadata. Otherwise the connection will fail during the reverse engineering process. (104037)
- In the mining editor, if you execute a mining flow against a DB2 V8.2 database, you might get some SQL exceptions that contain the string `SYSIBM.SQLCAMESSAGECCSID|PROCEDURE`. To fix the problem, change the JDBC driver for your database connection to the one shipped with DB2 V8.2 instead of the JDBC driver shipped with the DWE Design Studio. To make this change, update the path in the **Class**

location field in the Edit Connection screen. This problem was observed for DB2 V8.2 FP10 on Windows, but it might occur on other operating systems.

- The model building operators in the mining editor (associations, sequences, clusterer, and predictor) and in the mining and dataflow editor (cluster table, find deviations, find rules, and predict column) allow you to specify either a fixed value or a variable for the name of the model that is to be generated. The fixed value always has two parts, a prefix and the actual model name, separated by a dot. For variable values, the system does not enforce this two-part syntax, so that variable values are not required to include the dot. However, if you use a variable for a mining model name that does not have the dot, the execution of the mining flow fails. To prevent the error, always include a dot in the definition of variables used for mining model names. Additionally, for the cluster table, find deviations, find rules, and predict column operators, you must use the fixed prefix IDMMX and the model name must be in uppercase. (144721)

DWE Administration Console

Issues exposed in version 9.1.2:

- In the DWE Administration Console running on UNIX and Linux, after you complete the Advise Operation and then you try to execute the `advise_sql.sql` script using **DWE OLAP > OLAP Optimization > Run SQL Script**, the script does not execute, and no error message is displayed. For example, assuming that the installation directory is `/opt/IBM/dwe` and the server name is `server1`, take the following steps to fix the problem:
 1. Stop the WebSphere Application Server (WAS):

```
/opt/IBM/dwe/appServer/profiles/default/bin/stopServer.sh server1
```
 2. Remove the contents of the WAS temp directories:

```
/opt/IBM/dwe/appServer/profiles/default/wstemp  
/opt/IBM/dwe/appServer/profiles/default/temp/your_node_ID/server1/DWEAdminConsole
```
 3. Verify that DB2 is sourced into the environment.
 4. Restart the WAS:

```
/opt/IBM/dwe/appServer/profiles/default/bin/startServer.sh server1
```

Issues exposed in versions 9.1.0 and 9.1.1:

- Deployment of a data warehouse application fails under the following conditions:
 - In the Deploy Data Warehouse Applications page, the **Execute Deployment Code Units** option is set to **Yes**.
 - An SQL script that was included during deployment preparation in the Design Studio either did not reside in the current workspace or resided in a project that was not referenced by the current data warehouse project.

When you prepare new data warehouse applications for deployment, be sure that the required SQL scripts reside in the current workspace and in either the current project or a referenced project. (144528)

- In the DWE Administration Console, when you follow the wizard to create a database profile, and you choose to use a new security alias on the third step, you need to restart the WebSphere Application Server before you can successfully test the database connection. Because the J2C authentication aliases are loaded at the WebSphere Application Server startup, you must restart this server when you create a new authentication alias so that it can be loaded at that time. Then, you can successfully test the database connection. (144194, 144195)
- When you use the DWE Administration Console to enable a database for mining, and then use the mining command line to disable the database for mining, and you repeat this process multiple times, you may incur problems with the database application heap size. If you need to repeat the enable and disable processes multiple times, use the `idmenabledb` and `idmdisabledb` commands at the mining command line to avoid these problems. (144527)
- Archiving option for process logs: If you archive a process log and then re-run an instance that uses the same log file name, the log might be empty. To fix this problem, clean up the log directory

regularly by removing the `html.n` files. You can also change the log prefix before running new process instances to ensure that a new log is created. (22718, 23945, 24587)

- Console pages return error 500: Console pages sometimes return a null pointer exception (error 500). After this error occurs once, it tends to recur on other pages that you access. To solve the problem, restart WebSphere Application Server. (23886)
- Error messages about DB2 starting: When you access pages in the DWE Administration Console while DB2 is not started, you get an error stating that DB2 is not started; this is normal behavior. When you access pages a second time with DB2 not started, the error message does not display, although it should.
- Application names must not contain forward slash and backslash characters: During deployment, the name that you enter for a data warehouse application must not contain either backslashes (\) or forward slashes (/). These characters are interpreted as directory separators and cause the deployment process to fail. The error message that is currently returned in this case is misleading. (23685)
- Long path names can cause control flows to fail during execution because of operating system limits on the maximum length of a path. These limits apply to control flows that you run in both the DWE Design Studio and the DWE Administration Console. The length of a control flow path is influenced by elements such as the installation path, the control flow name, and the activity name. Using long names for any of these components can cause the maximum path length to be exceeded and the control flow to fail. Use good judgment when selecting installation paths and names. (24207)
- After successfully configuring WebSphere Application Server for the DWE Administration Console, if you change your password in the local OS for either the WebSphere administration user or the user that is used to authenticate to the databases, you also need to change your password in WebSphere. If global security is enabled, take the following steps to change the WebSphere administration user password:
 1. Login to the WebSphere administration console.
 2. Click **Global Security**.
 3. Under **User registries**, click **Local OS**.
 4. Update **Server user password** and click **Apply**.
 5. Save the WebSphere configuration.
 6. In the WebSphere installation location/`profiles/default/properties/soap.client.props` file, modify the password at this line: `com.ibm.SOAP.loginPassword`.

Whether or not global security is enabled, take the following steps to change the J2C Authentication alias used by the data sources:

1. Login to the WebSphere administration console.
2. Click **Security > Global Security**.
3. Under **Authentication**, click **JAAS Configuration**, then click **J2C Authentication Data**.
4. If there are multiple entries, choose the entry for the DWE Administration Console. For example, if your WAS hostname is `myhost`, click the entry named `myhostNode01Cell/etlAuthentication_myhostNode01_server1`.
5. Update your password and click **Apply**.
6. Save the WebSphere configuration.
7. Stop and start WebSphere.

Change the administrative account password on the DB2 server. For example, on Windows, you can change the services startup login password by taking these steps.

1. Click **Start > Settings > Control Panel > Services**.
2. For each of the following services, click the name of the service in the list, display the right mouse menu, click **Properties > Log On**, and then specify the new password:
 - DB2 - DB2-0
 - DB2 Remote Command Server

(24470)

- Single quotes in application descriptions and comments: When you deploy a data warehouse application or modify its properties, you can view and update the Description and Comment fields for the application. If the text in either of these fields contains single quotes ('), you might encounter SQL errors when you run processes that belong to the application. To avoid these errors, either remove the single quotes or use two single quotes instead. (24442)
- Console must be started *after* the DB2 instance that contains the runtime control database: If you start WebSphere Application Server and the Administration Console before you start the DB2 instance that contains the runtime control database (SQWCTRL), you cannot run data warehouse application processes. After starting DB2, restart the application server and console in order to manage and run your applications. (24447)
- Unsuccessful test connections for remote databases: Database connections that are required for data warehouse applications can be made either through the WebSphere Application Server connection pool or directly using a DB2 client. To connect to a remote database via the DB2 client, you must specify the authorization (user ID and password) to access that database when you define the data source in the console (Create or Manage Data Sources page). Remote databases that do not have this authorization defined return an error message when you use the Test Connection feature. You can ignore this message if the direct connection is not needed and you intend to access your remote databases exclusively through the WebSphere connection pool. (24423)
- DBCS characters not correctly displayed on the Manage Mining Models panel when importing a mining model: Here is a possible scenario:
 1. Expand the DWE Mining folder on the DWE Administration Console navigation panel and click **Model Management**.
 2. Import a new model with Japanese model name and description. The model is imported but the Japanese characters are not displayed correctly on the Manage Mining Models panel.To correct the problem, import the model from the Design Studio. Connect the Administration Console to the same database that contains the imported model. The DBCS should display correctly on the Manage Mining Models panel. (24336)
- Users who access the console from a web client can use multiple top-level browser windows (launched either from Program Files or from a desktop or quick launch icon) to administer the warehouse. However, the use of child windows within a top-level window (created by selecting **File -> New -> Window** from within an existing browser window) is not supported; the child windows share the session with the parent window and each other, which can cause inconsistent behavior. (24504)
- Depending on your browser security settings, you might be prompted to trust an unsigned Java applet when you use the mining model view in the DWE Administration Console. (22114)
- A process instance name can only be reused after it has been permanently deleted. Removing an instance from a single page view does not necessarily delete the instance. Deleting the same instance from both the monitoring page and the statistics page with **Retain Instance Statistics** set to **No** in the application properties page permanently deletes the instance. (24589)
- Data warehouse applications cannot be deployed to the Administration Console if the names of data flows and other objects that are part of the deployment package equal or exceed the following limits (number of characters). These limits are imposed by the column lengths that are predefined for tables in the control database; therefore, the limits apply only when applications are deployed.
 - Application: 128
 - Control flow: 220
 - Activity: 220
 - Data flow: 128
 - Mining flow: 128
 - Resource reference: 64 (database or machine resource name, as defined in a resource profile)
 - Variable ID: 128

- EPG: 255 (the relative location of the EPG file)
- EPGname: 192 (a subset of the EPG)

(24765)

- Do not use browser buttons to navigate pages in the DWE Administration Console. Use only the application buttons and links to navigate pages. If you use the browser back and forward buttons, you put the application in an undefined state and an error will occur. (94130)
- Before you enable a database for mining on a specific table space, consult your database administrator to ask which table space you can use or to request a new table space for your mining process. (143607)
- The tabs feature in the Firefox Web browser is not supported in this release. (144201)
- Before importing a large Cube Views metadata XML file that is greater than 1 MB, increase the Java Virtual Machine (JVM) heap size. Otherwise, you might encounter Error 500: Java out of memory from the WebSphere Administration Console. (144364)

To increase the JVM heap size, complete the following steps:

1. Copy the file `SQLLIB\MISC\db2mdapi.sql` to a new location and rename it.
2. Edit the new copy to increase the 1M default parameter size. For example, to increase the parameter size to 4M, change the lines as shown in the following example:


```
DB2INFO.MD_MESSAGE(IN request CLOB(4M),
  INOUT metadata CLOB(4M),
  OUT response CLOB(4M))
```
3. From the command line, connect to the database and execute this modified script: `db2 -<tvfcopy_of_db2mdaip.sql>`.
4. In the WebSphere Application Server, increase the JVM size.
 - a. Start the WebSphere Administration Console.
 - b. Select **Servers > Application Servers** and click the link for your server.
 - c. Select **Server Infrastructure > Java and Process Management > Process Definition > Java Virtual Machine**.
 - d. Enter values for the **Initial Heap Size** and **Maximum Heap Size**. (For example, enter 512 for the **Initial Heap Size** and enter 1024 for the **Maximum Heap Size**.)
 - e. Click **Apply**.
 - f. Click **Save**.
 - g. Restart WebSphere Application Server to apply your changes.
- Under **SQL Warehouse Tool**, on the **Resources > Manage Data Sources** window, when you test the connection to the control database, the test might fail and display an error message. The workaround is to restart WebSphere Application Server. (143583)
- If you install DB2 Alphablox, there is a defect that incorrectly loads JDBC files from DB2 8.2 into the DWE administration console. These back-level JDBC files cause errors in the OLAP functions. To correct this problem, delete the DB2 8.2 JDBC files and replace them with their DB2 9.1 equivalents. The following instructions assume that you accepted the default value for the installation directory.
 1. Stop WebSphere Application Server.
 2. Delete the following DB2 8.2 files from the `C:\Program Files\IBM\dwe\appServer\lib\ext` directory:
 - db2java.zip
 - db2jcc.jar
 - db2jcc_license_cu.jar
 3. Copy the DB2 V9.1 versions of these files from `C:\Program Files\IBM\dwe\SQLLIB\java` and paste them in `C:\Program Files\IBM\dwe\appServer\lib\ext`.
 4. Restart WebSphere Application Server. (145140)

- If you specify a remote database for the control database in the configuration tool, and then test the connection to it in the DWE Administration Console, an error message might result. If you can continue using the control database, then the error is incorrect and you can ignore it. (144239)
- On Windows, when you start the DWE administration console and create a database profile, you might see an error message about a failure in loading the T2 native library db2jcc2. To work around this error, restart WebSphere Application Server. (145017)

DWE Intelligent Miner

Issues exposed in versions 9.1.0 and 9.1.1:

- On AIX, if you are running DWE Intelligent Miner on a 64-bit DB2 instance with AIX ServicePack 5300-04-01, you might see error message SQL0444N when your user-defined functions (UDFs) are processed. In some situations, the service pack prevents DB2 from loading Intelligent Miner shared libraries. For example, you might see an error if you run the following command:

```
db2 "values IDMMX.DM_MiningData()..DM_defMiningData('OK')..DM_getTableName()"
```

The error message looks like this:

```
SQL0444N Routine "*ABLENAME" (specific name "SQL051005183118801") is implemented with code in library or path "...qllib/function/idmx", function "udf_getTableName" which cannot be accessed. Reason code: "9". SQLSTATE=42724
```

You can check whether the AIX ServicePack 5300-04-01 is installed on your system by running the following commands:

```
oslevel -r
oslevel -s
```

If the command `oslevel -r` displays 5300-04 and the command `oslevel -s` displays 5300-04-01, then ServicePack 5300-04-01 is installed.

The solution is to install APAR IY83580, which is available here:

www.ibm.com/support/docview.wss?uid=isg1IY83580

- In Miningblox, **GainschartBlox** doesn't work with Classification models if the WebSphere Application Server JVM is started in headless mode. You can solve this problem by disabling headless mode and restarting your server. On disadvantage to this solution is that it disables DB2 Alphablox if WebSphere Application Server is running on a UNIX system without a display.
- If you run Intelligent Miner Visualization as an applet, the parameter **noExport** is not detected for association and sequences models.
- Some Intelligent Miner table function like **DM_getFields**, **DM_getCorrelations** and **DM_getTreeRules** do not work properly for databases that are enabled for Version 8.2. of Intelligent Miner. For example, on Windows, the **DM_getFields** table function always returns empty strings as the values of the FIELDNAME column, and on Linux, the three functions fail. To solve the problem enable the database for Version 9.1.1 with the **idmenabledb** command. (CQIDM0002044 and CQIDM0002049)

Documentation Updates

The following information is provided as a supplement to the current DWE documentation. This information will be integrated into the Information Center in a future release.

DWE Design Studio

Issues exposed in versions 9.1.0 and 9.1.1:

User temporary table space for SQL execution databases: Be sure that each SQL execution database that you use in your data warehouse applications contains a user temporary table space. When data flows are executed, temporary tables are often created and this table space is required.

Table target operator: When you select the **Delete** option for the table target operator, you can apply a condition to the rows in the target table as well as a match condition for the incoming data. The Condition page has the following properties:

- **Input match condition:** For the SQL Delete operation, specify a condition that matches rows from the input data set with rows in the target table. For example: `PRODUCT.LINEID = INPUT_01.LINEID`
- **Delete condition:** For the SQL Delete operation, specify a condition that will be applied to the target table. Only the rows that meet both this condition and the input match condition will be deleted. For example: `PRODUCT.PRODUCTID >100`

Bulk load target operator: When you run a data flow that uses the bulk load target operator, descendent foreign key tables and materialized query tables can be left in the check pending no access state. You have to run SET INTEGRITY statements manually on these descendent tables to bring them back to a normal state. A SET INTEGRITY statement is run only for the table target itself.

Restrictions on data flows that can be deployed as DataStage executables: Data flows that generate Java code units cannot be deployed as DataStage executables. This restriction applies to data flows that contain the following operators:

- Table source and table target operators with the **Location** field set to **Remote database**.
- Table target operators with the **SQL operation** field set to **Insert** and the **Commit interval** field not set to 0.

SQL replication source operator: If the Apply server runs prematurely before the Capture program has fully completed committing the source changes to the CD tables, then the changes in the source table might not be consistent with the rows applied in the replication target table (especially, if partial changes are captured).

In such a case, start the Capture program externally (usually a daemon process). This causes the Capture program to run asynchronously with the SQW process that executes the Apply program. Timing inconsistencies, especially due to system load, could occur when data is captured and applied at the same time. This is possible during both Design Studio executions and scheduled process executions in the Administration Console. You should account for such timing inconsistencies while scheduling your application processes in the Administration Console.

You are recommended to schedule the execution of the Apply process after you are certain that the source table changes are completely captured or you must schedule the process execution multiple times over a period to eventually apply all the changes to the replication target tables.

Default image format for saved flows: The documentation states that the default image format when you save a flow as an image is JPEG. In Version 9.1.2, the default format is GIF.

DB2 Data Warehouse Edition, Version 9.1.2 (3/28/2007)

Licensed Material - Program Property of IBM

IBM, DB2, DB2 Universal Database, AIX, Cube Views™, DataStage, Intelligent Miner, Rational, and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both. Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

© Copyright International Business Machines Corporation 2006. All rights reserved.

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