What’s new in IBM DataQuant Fix Pack 6

We have continuously updated DataQuant and added enhancements to the package in fix packs. It may be helpful to review these enhancements as sometimes such knowledge is just ‘assumed’ to be commonly known. Now here is a comprehensive update that will enlighten you to the many exciting new features we have added.

Virtual Data Sources – a new metadata layer

In prior versions, users developing queries, reports and dashboards were required to understand the underlying physical tables and views from which their queries drew data. With the introduction of virtual data sources in fix pack 6, administrators can now optionally shield their users from the complexities of the underlying database structures, and provide content developers with a simplified data model against which content can be developed.

Virtual data sources work by introducing a metadata layer which mediates between an administrator-defined, virtual data source and the underlying data sources that contain the physical tables and views. This approach provides the following key benefits:

1. End users are shielded from the complexities of the underlying data model. Obscure column names and table names can be replaced with human-readable variants and complex table joins can be represented by a single virtual table, much like a database view.

2. Administrators can define multiple virtual tables in a single virtual database, each of which draws data from one or more tables within differing data sources. For example, one could add three tables from DB2 and two tables from Informix. To end users, a given virtual data source acts as a single database, allowing end users to author queries against all tables contained within it, despite their underlying sources residing in differing data sources.

3. End users are insulated from database schema changes. Database administrators can readily make changes to the underlying database schemas without affecting existing queries, reports and dashboards, since the virtual data source schema remains unchanged.

Job Scheduler – automate repetitive work

IBM DataQuant introduces the ability to schedule a DataQuant procedure object to be executed on a periodic, recurring basis or at a specific date and time. A DataQuant procedure can perform one or more sequential steps, such as run queries, apply report templates, export reports to PDF files, send reports via email etc.

Jobs can be scheduled to run at either of the following locations:

1. Local workstation: DataQuant automatically creates a job with the operating system scheduler, under the profile of the logged on user. All execution takes place at the end user's workstation.

2. DataQuant for WebSphere: DataQuant schedules a job for execution by DataQuant for WebSphere's new built-in job scheduler. In this mode, the job is executed by DataQuant for WebSphere itself, fully independent of the workstation that scheduled it. Using this approach, multiple users may schedule work to be performed on an ad-hoc or periodic basis, without the need to leave their workstations powered up and logged on.

Support for Multi-Language Dashboards – go global!

IBM DataQuant's interactive visual dashboards provide an excellent means of disseminating your business data and reports across the enterprise. For global companies, this often means deploying solutions across differing language regions. While IBM DataQuant is translated into over 20 languages, any text that you add to your dashboards remains in its original form, regardless of region.

IBM DataQuant now includes the ability to author dashboards that will present localized text that matches the region of the end user. For example, you can now author a dashboard in German that also contains your text in French. In this example, German users viewing your dashboard see German text, whereas French users accessing the very same solution see your French text. This new feature allows you to define as many 'string tables' as you wish within a given dashboard, each of which can be assigned to a given regional language.

Support for Flash-based Dashboards – increase the sizzle

In prior versions, DataQuant dashboards could be deployed using pure HTML, PDF or optionally via a lightweight browser Applet. IBM DataQuant now introduces the ability to deploy dashboards using Adobe's Flash player. The Flash-based option provides an end user experience that is more comparable to a rich application since the dashboard is *assembled* by the workstation and many dashboard interactions can take place without the need to communicate back to the web server.

Introduction of backward-compatible Win32 APIs – install protect

IBM DataQuant offers three forms of application programming interfaces (APIs), allowing DataQuant functionality to be optionally embedded within custom applications and third-party software. These include Java APIs, Web Service APIs and Command Library (command line) APIs.

This new feature introduces a set of Win32 APIs that provide access to DataQuant functionality using a set of functions that are compatible with the former QMF for Windows APIs. This allows in-house applications that were developed using the QMF Win32-based APIs to operate with the newer Java-based DataQuant solution.

Data Expiration Schedules in Reports and Dashboards – more data access control

In prior releases, DataQuant's visual reports and dashboards retrieved data from the database when the solution is accessed. For example, a sales report returns sales data that is current at the time of opening the report.

DataQuant's new *data expiration schedules* allows dashboard and report developers to define the point at which new data should be fetched from the database. Expiration schedules can be defined in terms of a time span (such as so many seconds, minutes, hours, days etc.) or as a regular interval (such as every day at noon or every first Monday of each month etc.). In addition, each query in a dashboard or report can be assigned a distinct expiration schedule.

This new feature can be utilized in many ways, such as:

- Defining a sales report that always displays data for the current month. The data is fetched once only, after being accessed for the first time after the new data is available. Subsequent access to the report is satisfied by data held in the report's data cache. Copies of the report can also be saved against a given month, creating a repository of report snapshots on a monthly basis.

- Defining a dashboard that refreshes each query according to the underlying validity of the data. For example, book authors never expire whereas transactional records are valid for so many minutes. By pooling the data caches at the web application server, concurrent user access to dashboards results in virtually no incremental loading on the underlying data sources that power them.

Enhanced IBM DataQuant for WebSphere User Interface and Capabilities – thin is in!

IBM DataQuant's Eclipse-based user interface offers a productive and increasingly familiar environment for both technical and non-technical users alike. In prior versions, this interface was limited to rich client users. DataQuant for WebSphere's user interface offered a subset of workstation features but using a distinct HTML interface.

IBM DataQuant for WebSphere now provides web users with a user interface that is virtually identical to the workstation product. This retains the benefits of a pure-HTML solution without losing the productivity of a desktop application. Application toolbars, context menus, dialogs and views remain common between both products, allowing for the seamless transition between each of the solutions.

New 'Personal' View – better user placement and organization

IBM DataQuant now includes an optional 'Personal' view that is docked above the workspace or repository explorer in both the workstation and web products. This new view offers the following facilities:

- A 'Recently Used' folder item that lists and provides direct access to tables, queries, reports, procedures and dashboards recently accessed by the end user.
- A 'Favorites' folder within which end users can collect their 'favorite' items (tables, queries, reports, dashboards etc.). Items are added by dragging and dropping or by right-clicking on objects and selecting 'Add to favorites'.
- A 'Startup' folder that allows users to define items that should be opened whenever the application is launched. Items added to this folder (using the same mechanism as above) are automatically opened and displayed when the user logs into the application.

What's new in DB2 QMF Fix Pack 10

All of the above features also apply to QMF for Workstation, with the exception of the Flash-based dashboards and Multi-lingual dashboards.

For more information on IBM's DataQuant please go to:
http://www-01.ibm.com/software/data/db2imstools/db2tools/dataquant/

For more information on QMF please go to:
http://www-01.ibm.com/software/data/qmf/