

**(Intro)**

Today we'll learn how Integrated Data Management solutions from IBM helped make one DBA's day a little better and a lot more productive.

**(Jackie typical day slide)**

Jackie is a typical DBA with too much to do and not enough time to do it. Her tasks usually fall into one of the following areas:

- **Maintenance**, including space and resource management, and doing proactive performance analysis.
- **Application management**, including rapidly responding to requests to create new structures or change existing structures, grant authorizations, and provide good quality test data.
- **Availability**, meaning she must devise and implement an appropriate backup strategy and is responsible for recovering data.
- Finally, **performance and tuning**, which means she must react with high speed to any problems in production environments, diagnose response time issues in complex application environments, and tune the system to meet service level agreements.

**(Tools she'll use today)**

Jackie is lucky in that she has access to solutions to help her do her job better and faster. Today, she'll be using the following tools from the IBM Integrated Data Management portfolio:

- DB2 Performance Expert and Extended Insight Feature to help with diagnosing complex performance problems and for proactive monitoring.
- Optim Database Administrator for simplifying backup and for managing application changes. She'll use Optim High Performance Unload to help a colleague find a solution to unload and migrate data quickly.
- And she will use Optim Test Data Management to create a right-sized test database that conforms to her company's requirements for managing private data.

**(Performance Expert)**

**The following happens between 2 AM and 3 AM**

Being a DBA is a critical role in the organization and can sometimes mean being paged at all hours. This happens to Jackie today when her cell phone wakes her with an email alert from DB2 Performance Expert indicating that

a problem threshold has been exceeded on one of the System Test systems. She set up this alert because this application is set to go into production soon, and she needs to be alerted to any potential problem before that happens.

She views the DB2 Performance Expert System Overview screen, which acts as a “dashboard” to see high level performance information about all DB2 systems she administers. She notices that there are an unusually high number of sorts per minute on the test server. She drills in further to see the key performance indicators for the Great Outdoors sales database.

The guidance provided by Performance Expert lets her know that one of the KPIs is too high and that this might be a problem in which an index can help.

At this point Jackie still doesn't know which application is experiencing the problem, so she switches to the application summary view, and refreshes the performance data.

Then, she can sort by elapsed time to bring the application with the longest elapsed times quickly to the top. She then selects **application details** to see detailed metrics for that application.

First, on the buffer pool information details, she notices that there are very few or no index reads when compared to how much data is being read

She continues her investigation by looking at the details of the SQL activities and sees a graph indicating a high ratio of rows read to rows selected. This indicates the table is being scanned before returning the rows needed to satisfy the query.

To confirm her hypothesis that an index is not being used, Jackie drills down into the SQL Statement and Package information. From this view, she can see the most recent operation and the actual statement that is running. To see the access plan for this SQL statement, Jackie launches DB2 Visual Explain. She can see that two table scans are being used instead of indexes, which surprises her since they didn't have this behavior on another QA system.

She clicks directly on the Visual Explain graph to investigate one of the table scans in more detail. She sees the list of columns that are retrieved by

this table scan, which should be considered to be included in the index or indexes that she needs to create on the SVT system to solve this urgent problem and ensure the performance is brought back into an acceptable range.

**(Performance Expert Extended Insight)**

**The following takes place between 8 AM and 9 AM.**

Hoping for a more routine start to the day, Jackie brings her coffee to her cubicle and does what she does most mornings when she comes to work – check out the Performance Expert System Overview which is an enhance with the Extended Insight Feature for end-to-end database monitoring. Unfortunately, there is red showing, indicating that a problem threshold has been triggered.

She follows the red indicator all the way through to the view that contains the two workload clusters that are having problems. Workload clusters are simply different ways of viewing performance information.

Jackie scrolls down to get more details on which application and client user id is responsible for tripping the threshold. She discovers that the business report application has exceeded its problem threshold.

Now that she knows which application is causing the problem, she continues her analysis by investigating which Client User Id is experiencing the slowdown.

Jackie sorts the client user IDs and discovers that the BI Report user ID has exceeded its problem threshold. A quick glance at the end-to-end graph for the client user ID quickly enables Jackie to find out when the response time spiked above the red problem threshold and the yellow warning threshold. It's becoming clear that the problem is occurring for some reason because of the business reports being run on the system. Now she must determine where in the software stack the problem is occurring to see if this is a database problem or something else. She continues her investigation by double-clicking on the BI Report user ID to drill down and view the end-to-end details.

Jackie changes the data to focus on the activity that has occurred within the last 15 minutes. She can see the top 3 SQL statements that were issued from the Report application during that time. Jackie selects the SQL statement

with the longest end-to-end response time and drills into the detailed statement information. The pie chart shows the distribution of factors responsible for the end to end response time for this statement, and she can see that 99% of the time is in the data server indicating that she has more work to do to figure out what's going on.

Jackie launches the DB2 Visual Explain to view the access path for this SQL statement. She sees the problem seems to be a nested loop join that is not using an index and causing the cost of that join to jump dramatically.

Based on the data she now has, Jackie calls the business analyst responsible for the report to suggest that they should work together on rewriting the query.

**(Summary of PE & PEEI features & benefits)**

By using DB2 Performance Expert with its Extended Insight Feature, DBAs can significantly reduce the time they spend diagnosing problems in multi-tier environments for Java applications. For Java and non-Java workloads alike, Performance Expert helps DBAs keep on top of emerging issues to avoid critical outages or slowdowns with trend analysis and customizable alerts.

**(Optim Database Administrator for Database Administration Tasks)  
(Flat View)**

**The following takes place between 9AM and 10 AM.**

Jackie needs to prepare for an upcoming performance and stress test of a new application. She needs to verify that there is enough sales data in the system test environment to adequately stress the application. She can quickly do this using Optim Database Administrator.

From the Data Source Explorer, Jackie toggles to a 'flat view' of her objects. This flat view returns all objects of a particular type, such as tables. She clicks on the Table folder of the test database, and then sorts by Schema to find the test tables and then by Row count to quickly locate any tables that might not have enough data.

**(Database Connection Working Set)**

Jackie normally manages dozens of databases, but her focus right now is being able to respond rapidly to changes required for a new application. She needs a way to quickly get to the development and QA databases. She decides to set up a Working Set of database connections in Optim Database

Administrator. This working set is in essence a view of only the database connections that she defines.

To do this, Jackie names a new working set. Then she simply adds the desired database connections to the working set. After the working set is created, only the database connections she needs for this project are shown. She can create other working sets as needed to include other connections she may need for particular projects. Jackie can now respond with greater agility to the dynamic environment in which she works.

**(Utility Commands in Data Source Explorer)**

Jackie is a strong believer in a good backup and recovery strategy. Now is the time to back up the SVT system, which she can do directly from the Data Source Explorer. As with any Eclipse-based workbench, she can right click on the database to see all operations that can be performed. She selects Backup, which takes her directly to the task assistant which helps her perform the operation.

As a sanity check, she previews the list of commands needed for the Backup operation before running the backup. (If needed, she can edit the commands before running them.) The status shows that the script ran successfully.

**(Optim Database Administrator for Database Change Management & Deployment)  
The following takes place between 10 AM and 11 AM.**

The new application project continues to move along at a rapid pace. Jackie needs to respond to an urgent request from QA to migrate two new tables and their data from the development system to QA for testing. This is easy to do using Optim Database Administrator's copy and paste capabilities.

First, Jackie uses the flat view to find the new PARTNER ACTIVITY and CONTACT tables in the development database. She right clicks on the tables and chooses to Copy. Next, she finds the target Test database, and pastes the copied data. Because QA needs the data for testing, she copies both the data and the database objects. This takes her to the Change Management Script Editor where she can preview the list of commands for the migration, including the import of the data. She also decides to get a Summary of Changes Report to see a nicely formatted summary of the changes indicating that two new tables with primary keys and new indexes will be created. If she wants it, she can also send this report to the QA team to ensure that the proposed change meets their requirements. For now,

Jackie runs the change script to create and populate the new tables, and to create the indexes. The result indicates that the Change script has been run successfully. Note that if anything goes wrong, Jackie can click on the Undo command to back out the changes.

Next, Jackie validates the changes by finding the partner activity and contact tables in the test system and getting the contents of the partner activity table. The QA team is impressed at how quickly she was able to respond to the change request.

**(Summary of Optim Database Administrator's features and benefits)**

Optim Database Administrator is the desktop tool every DBA needs. It includes utilities and commands for day-to-day maintenance, and also helps DBAs manage database changes from simple to complex. It automates complex structural changes while preserving data, privileges, dependent objects, and application bindings. It reports on the impact of proposed changes to identify dependencies and reduce the risk of application outage and provides an easy way to undo changes. Finally, it documents a complete history of the changes for auditing purposes. With Optim Database Administrator, DBAs can react with greater agility to a dynamically changing environment.

**(Optim Test Data Management)**

**The following takes place between 3 PM and 4 PM.**

Jackie receives another urgent request from QA who wants good quality customer and credit card test data to validate new application features. They would like to use a subset of existing production data to do this. To avoid data privacy compliance violations, Jackie must de-identify the data before exposing it to QA. Luckily, Optim Test Data Management enables her to create a viable subset of data and the Data Privacy Option lets her privatize the data such that valid testing can be performed.

Here are the steps Jackie will follow using Optim Test Data Management with the Data Privacy Option:

- Extract a valid subset of data from production into a compressed extract file
- De-identify and insert the data from the extract file to the test database
- Compare production and test data to verify that the credit card data is de-identified

Jackie generates a new Extract Request. She defines an access definition, which describes the set of related data to perform an operation on – in this case, to subset. These access definition files are handy because Jackie can edit or reuse them for other purposes such as subsetting to another development system or for archiving.

On the Access Definition Editor, she specifies the tables to be extracted. She wants to extract from the production system so she selects the production qualifier. Now she can start selecting tables to extract. She chooses the Credit card table as her start table. She decides she wants to add child tables related to credit card as well. Optim can find up to any level of parent and child relationships automatically.

For this test environment, Jackie needs just one level of child relationships, and you can see here that the customer order child table is added to the list of tables to be extracted.

Jackie has many options to specify subsets of data. In this case, Jackie just wants to get the first 100 rows from the credit card table. Rows from the child orders table will be automatically determined by the rows extracted from the parent Credit Card table.

Jackie now runs the extract and can see from the Extract Process Report which objects and how many rows were extracted.

Next, Jackie creates an Insert Request to populate data from the extract file to the QA database with data privacy applied.

She defines a mapping of the table columns from source to target, which also is how she can specify a data privacy algorithms and/ or transformation functions to privatize the data. She selects the credit card table and indicates that she wants to create a column mapping for that table. She wants to change the credit card primary account number to use a fictional, de-identified credit card number using one of the built-in functions in the Optim Data Privacy Transformation Library. “Transformation” is shown on the right column indicating that the value for this column will be transformed. Other columns will have the same data as the source.

Now that mappings and data privacy are defined, Jackie starts the process to insert data into the test database.

Finally, to verify that the data was privatized correctly, Jackie uses Optim's Compare function to compare data extracted in the extract file and data inserted in the Test database. The comparison feature shows that entries in Customer Order tables in source and target are the same (in other words, there was no data privacy applied), and that entries in Credit Card tables are different because data privacy was applied. To drill down on the differences, Jackie double clicks on these tables. Each pair of the tables is shown with the difference columns in bold

**(Summary of Optim TDM features and benefits)**

IBM Optim Test Data Management Solution enables you to create realistic, right-sized test databases, edit data to force error and boundary conditions, and compare baseline data against successive test results to identify application errors that would otherwise go undetected. In addition, the Data Privacy Option makes it easy to comply with privacy requirements for test data. By streamlining test data management, companies can improve application reliability, speed iterative testing tasks, and lower development and testing costs throughout the application lifecycle.

**(Optim High Performance Unload)**

**The following happens between 4 PM and 5 PM**

Erik stops by Jackie's office because he has a problem. He only has a few hours during Sunday's downtime to unload a large quantity of data from the production database. He doesn't think he can meet this requirement using DB2 Export and wants to know if she has any advice. She tells Erik that he should try using Optim High Performance Unload, a high-speed utility for unloading, extracting, and repartitioning data in DB2, which is often up to 10 times faster than using the DB2 export command alone. It can perform faster by directly reading from the database files, and it is highly threaded to exploit the hardware to perform the unloads in minimum elapsed time through parallel processing.

She is able to show him how fast it is directly from Optim Database Administrator, because of built-in integration between the two products. It's important to know that Optim High Performance Unload can also be invoked from a DB2 command line.

She chooses the ORDER\_DETAILS table to unload, then invokes Optim High Performance Unload. She enters file destination for the unloaded data. They preview the command, and then run it. She shows Erik that the utility unloaded over 4 hundred thousand rows in about 4 seconds. Erik is glad he stopped by – he thinks Optim High Performance Unload will help him meet his requirements to minimize impact on production and to repartition data for the warehouse.

**(Summary)**

Optim High Performance Unload helps DBAs meet service level agreements with high speed unload capability. It includes ability to repartition data and can filter columns and rows with simple predicates and even do some basic transformation on the data as it is being unloaded. It can help speed backup and recovery operations by enabling unload and recovery of single tables from a full database backup. These capabilities ensure that your data management team more flexibility in responding to a variety of outages.

**(Summary – slide 1)**

Today we've seen how Jackie was able to do the following:

- Rapid diagnosis of a performance problem
- Proactive monitoring to prevent future emergencies
- Database maintenance tasks that ensure speedy recovery from outages
- Support application requirements for data and object migration and with rapid creation of high quality test data that is compliant with data privacy requirements
- Fast unload of large amounts of data to support data migration while reducing the impact on production systems.

**(Summary – slide 2)**

All with the assistance of IBM integrated data management solutions.

**(Resources)**

Try Optim Database Administrator today by clicking on the download link. To find more information about Integrated Data Management solutions, visit our web site, and join the community.