DB2 9.7 – a technical overview

By Raj Mahal
Consultant & DB2 Expert, Triton Consulting

September 2009

DB2 gets more bite with COBRA

DB2 has been making huge strides over the past few years. The previous version , DB2 9.5 had made significant changes from it's predecessor, DB2 8.2 and DB2 9. The latest edition, DB2 9.7, aka COBRA, has incorporated an abundance of new functionality, which also offers much better compatibility with Oracle, also providing a much smoother migration from Oracle environments to DB2 environments, if the user sees this as a feasible option for their organisation.

Previously the work involved to move to the cheaper DB2 platform was considered to be weeks, if not months, this can now be achieved with relative ease and within an acceptable timescale.

Significant changes made to the DB2 system for this release incorporated some administration, database setup, application development, CLP and system command changes. Of these, the standout features include:

- extended security,
- database partitioning - available only as part of the data warehouse products
- changes in database manager parameters and registry variables
- automatic creation of data compression dictionary
- inclusion of logs into backup operations
- identifier length limits increased
- simultaneously backup of multiple databases and command changes for db2mtrk/db2ckmig/db2audit/db2pd.

Why has this change taken place?

As well as being the cheaper licensing option, there are changes that have been made and the introduction of new features that natively support Oracle features. To activate some of these features, a registry setting must be altered, as the features are under the registries control. The registry variable is ***DB2_COMPATIBILITY_VECTOR*** and is set to a hex string, this allows you to turn on the various individual features or all of them. Amongst these are;

- **Concurrency** – this is a key concept for databases. DB2 in previous versions had a default setting of cursor stability, which provided an adequate level of concurrency. However, as the number of transactions have increased significantly, the chances of locks and lock waits have also increased. New developments within DB2 version 9.7 means readers and writers don't block one another. The new default behaviour for DB2 9.7 is the currently committed isolation level, which replicates the Oracle default. Previously if a reader came across an uncommitted row, the cursor stability isolation level would mean the reader cannot access an uncommitted change. With the new isolation level the reader is able to see the currently committed value for the uncommitted row that is returned. The inner workings of such a change incorporate accessing the log buffer. If a row is being updated by a transaction that row will be locked and the transaction is placed in the log buffer whilst it is still uncommitted. If another transaction is executed against the same row, the transaction will find the row locked (cursor stability). However, changes made in 9.7 mean that the transaction finds the row locked and then accesses the log buffer to find the current value for the committed data. The new functionality can be enabled using the database configuration parameter **cur_commit**.
- **SQL** - the language used by most if not all relational databases, the language may seem the same to the casual user, however there can be subtle and vast differences from one platform to the next. The changes will be crucial to application developers as it will mean queries can be written that run on both Oracle and on DB2, a selection of SQL, which is compatible with Oracle

include, support for CONNECT BY recursion, NEXTVAL/CURRVAL have increased functionality so that they are more compatible with Oracle, MINUS is used in DB2 to represent Oracle's version of the EXCEPT keyword in DB2, the TRUNCATE statement allows for all the rows to be deleted within a table very quickly as there is no logging involved, ROWNUM limiting the amount of rows returned by a query. Some of these features do not require the registry variable to be set, NVL relates to COALESCE, DECODE is very similar to the CASE expression but has one or two differences, LEAST and GREATEST are similar to MIN and MAX. These are just some of the compatible features, there are many more.

- DB2 with its built in native compiler is able to compile PL/SQL statements, whether they are existing scripts, package definitions or single statements. Once processed the object is created in the database. Commands can be run either in the DB2 command window or using the command line processor, which is now SQL+ compatible. The SQL+ compatible processor is now more interactive supporting dynamic SQL and scripts. The support for this feature can be enabled using the ***DB2_COMPATIBILITY_VECTOR*** registry variable.
- DB2 is compatible with the Oracle data dictionary concept using views. Within Oracle the data dictionary is a repository for the databases metadata. The views used within DB2 for the data dictionary are based on the DB2 system catalog views.
- Stored procedures and the compatibility between various RDBMS is an important aspect for any organization, stored procedures are used by many if not all organizations. DB2 has in this version used the Oracle terminology, to enable any developer who is able to code in Oracle to code in DB2. This feature alone provides excellent benefits to organizations, as they mostly have many RDBMS and require code to be written for all.
- Online schema changes

Prior to the release of DB2 v9.7, trying to move to a version of DB2 would have required a huge amount of time and effort. However, with the new version of DB2 9.7 the compatibility feature allows for a much smoother transition and a lot quicker implementation for applications going from Oracle to DB2. With DB2 supporting around 95% of Oracle functionality. Given that the state of the current economic stability around the world, the benefits of DB2 and its lower CPU costs, storage costs and compatibility with Oracle, there is an argument for Oracle houses to move onto DB2.

Apart from native Oracle support, DB2 version 9.7 has incorporated increased functionality, key features include;

- Reduction in storage costs and storage administration
- SQL and performance improvements
- Enhancing workload management
- Enhanced security features
- Manageability Enhancements
- XML
- Statement concentrator

Prior to the release of DB2 v9.7, trying to move to a version of DB2 would have required a huge amount of time and effort. allows for a much smoother transition

<u>Reduction In Storage Costs And Administration</u>

IT budgets are used most by the cost of storage. The cost of storage is escalating, amassing larger proportions of the IT budget. As hardware costs are increased, the focus is moving towards using less storage within DB2. Less storage requires less management, a reduced use of power and a reduced amount of cooling for database management. Compression began in DB2 9.1 with table compression, It has now evolved in DB2 9.7; compression searches for repeating patterns and builds a compression dictionary. It then replaces the pattern with a 12-bit symbol. The symbols are stored within the

dictionary for a faster lookup. Other vendors look for repeating patterns in a single page or block, whereas DB2 will look at the whole table.

Compression for indexes is also yielding great results with faster performance and less storage. Temporary table compression is an excellent feature for large data warehouses, where there are lots of scans, order by's and sorts. As a temporary table grows and a threshold is reached the temporary table will be compressed. LOB in-lining can take a LOB and store within the row rather than in a separate table space. Now that they are stored within the row they are then eligible for table compression. With XML data compression the previous release allowed for in-lined XML, but now you can compress the XML data, using the repeating patterns method and the 12-bit method. This allows for an increased level of storage and less use of disk. If a table already has compression turned on and contains xml data, that data will be compressed automatically.

There will be a reduced amount of space required for backups as less pages are used, which also results in faster restores of the database environments. Index compression is completed automatically if table compression has been turned on. Index compression involves using the RID list for compression, so instead of using 6 bytes for each RID, you can use potentially 1 byte for each RID. Key values can also be compressed. By implementing index compression, huge improvements can be achieved for both updates and inserts.

There are also potentially huge savings when compressing temporary tables.

Tablespace remapping helps to reclaim space from tablespaces. If a new tablespace has been created, remapping will allow you to move the table around the tablespace. There may be a situation where some free space within a tablespace cannot be used for a new table. A new command has been added where you can alter the tablespace and specifically the high water mark. Once the high water mark has been moved then a new or old table can make use of the free space. This also benefits the backups and restores as you will be using less pages, which results in faster backups and restores.

Partitioning

Previously you could only create a global index for a partition within v9.1 and v9.5. In DB2 v9.7 you can now create local indexes that only point to an individual partition. It is also easier to roll-in tables now, such as a month table, without any need to carry out maintenance. So less resource is used. For example, if you create a table for June, add data, create the index, run the attach command, db2 will attach the table and index. Making it much easier to roll-in data and also detach data, rather than using global indexes, maintenance becomes a whole lot easier.

Runstats has also been enhanced to collect stats at table level as well as the partition level.

SQL and Performance improvements

Significant steps have been taken to deliver performance improvements for this version of DB2 to sustain and increase its position in the market as a leader in database technologies.

One of the improvements allows the user to make the most of access plan reuse, where, for static sql statements the access plan will either remain the same or will be very close to the query execution prior to this run. This allows the user to be confident that performance will be near expectations even when there have been version upgrades or statistical updates. In order to use such functionality apply the **ALTER PACKAGE** statement or apply the **APREUSE** option for BIND, REBIND and PRECOMPILE

commands. In relation to the former ALTER PACKAGE statement, you can now use this statement to relate optimisation profiles with a package, without having to execute a BIND.

This is just one of the improvements that have been made, further enhancements include;

Scalar Functions: ADD_MONTHS, DAYNAME, EXTRACT, LAST_DAY, MONTHS_BETWEEN, MONTHNAME, NEXT_DAY, TRUNC_TIMESTAMP and more

Created Global Temporary Tables: Created global temporary tables are now defined within the catalogs, users are able to reference them like other tables, CGTT are also under authorisation control of the user, you can create views as well as indexes and triggers on and LOBS.

Truncate Table: Rather than having to delete all rows, which can be very slow, or running a load using an empty load file. The new functionality in DB2 allows you to run the Truncate Table command, which is a lot faster and much tidier. There will also be no logging involved and no chance to rollback.

Statement Concentrator: There are a large number of statements that can be executed within OLTP environments and quite a few of these have very similar syntax and the only real difference between them is the literal value. Prior to DB2 9.7, the similar statements would require re-compiling individually and access plans generated for each statement, however generating plans for each statement is costly. With DB2 9.7 all similar statements will use existing access plans and not require any re-compiling for each and every statement. The statement concentrator will look to avoid the differences in literals within each statement and instead concentrate on the actual similarity between statements and then opt to use existing access plans, where possible.

Scan Sharing: scans begin at the first row and continue until the last row. In 9.7, you may have a million rows to scan, user 1 scans first row through to 1 million. If user 2 comes in and runs same scan then user 2 can jump in at 500,000th row, where user 1 had reached. So user 2 is piggy backing on user 1 and only using 1 i/o, then it goes back to the rows it missed.

<u>Enhanced Workload Manager</u>

Workload management is the concept of managing a diverse set of tasks, which can range from millions of inserts into the database environment, backups and restores of the database, reorganisation of specific tables within a database or even hundreds of reports being generated relating to sales achieved in countries. All of this work places huge amounts of stress on a database and its server, at this point workload management is vital to the sustainability of the database and the utilities running against the database. The management of a workload often requires working towards a set of goals, which maybe SLA based. Begin with identifying the work to be completed and whilst managing the work, assign the appropriate resources. Further monitoring is required, followed by evaluating any problems, all with the benefits of accessing real-time information.

DB2 version 9.7 has again delivered a solution with some enhanced features and some new ones to boot.

- More often than not a user will be aware of the kind of work being undertaken and the amount of management and monitoring required. However, at times it is difficult to anticipate the amount of work or management that will be required and at these times, some thresholds need to be in place. If these thresholds are breached then there needs to be something in place to manage these breaches. Some existing thresholds have been improved and a few new thresholds have been defined for 9.7, which include,

ESTIMATEDSQLCOST, SQLROWSRETURNED, ACTIVITYTOTALTIME, SQLTEMPSPACE, SQLROWSREAD, CPUTIME, AGGSQLTEMPSPACE, CPUTIME and SQLROWSREAD.

- The user can now allocate wildcards for similar workload groupings when defining a workload. Using the asterisk can be beneficial to highlight a number of jobs with the same application name. There is also the capability to use the wildcard with other workload attributes other than the application name.

- Integration between the DB2 workload manager and the LINUX workload manager.

- Thresholds that are based on time intervals have now been adjusted so that the minimum time interval has been reduced from 5 minutes to 1 minute. Specifically for ACTIVITYTOTALTIME and CONNECTIONIDLETIME.

- Changes made in the way longer running activities are managed. Within version 9.7 the priority for longer running activities have been lowered so that the shorter running activities can have some performance improvements. This is put into practice by changing the service class levels for the running activities. New thresholds available to use with this functionality, CPUTIMEINSC and SQLROWSREADINSC.

Enhanced Security Features

An increased level of security is required outside and inside database environments, especially now that databases and online businesses are required to run 24x7. DB2 9.7 enhances the existing security features, with an emphasis placed upon DB2 roles. These changes entail different attributes for different users, the database administrator can no longer have access to the data and can only maintain the database. The security administrator (SECADM) will look after security and auditing and can grant specific roles to groups. These new authorities will allow the users to only have what it is they require, such as WLMADM for workload management, SQLADM for sql tuning, EXPLAIN for the explain facility and so forth. The system administrator (SYSADM) is not implicitly given the DBADM authority and will not have the authority to grant, except tablespace, privileges. The security administrator can provide DBADM authority if required. On a smaller scale the SYSMON can now issue LIST commands and also make use of db2mtrk command.

Further enhancements include advanced encryption for user id's and passwords using the AES algorithm, as well as support for secure socket layer and transport layer security, which improves data security.

Manageability Enhancements

Further manageability developments in DB2 9.7 now allow for less involvement for the DBA. More and more tasks that previously required the intervention of the DBA are now automated or moving towards becoming automated. This results in the DBA having more time available to work on other issues or concerns within DB2. Some of the improvements that have been made are listed below:

Automatic Storage: this feature has been part of DB2 for a while, but now a DBA is able to convert existing databases and DMS tablespaces to use automatic storage. Using the ALTER DATABASE command, or for DMS tablespaces, using the ALTER TABLESPACE command or a redirected restore. Previously, you could only use automatic storage with new databases.

Storage Paths: as of version 9.7 the DBA is able to make use of any extra storage that has been added, there is also the option to drop storage paths if required. These options are available for databases with automatic storage. Changes can be made using the ALTER DATABASE with DROP STORAGE ON and ALTER TABLESPACE with REBALANCE commands. There is also new functionality that will allow the DBA to monitor the storage paths, available through the snapshot monitor.

Multidimensional Clustering: reclaiming extents from MDC tables that are not in use can be carried out through the use of a REORG. Whilst the REORG is running, access to the tables can be controlled, no access, read access and write access. The space when reclaimed can then be used by any database object, rather than just the MDC table.

Table Data Online Move: a new stored procedure can now move table data into a new table, which has the same name. The data will still be available to users for select, insert, update and delete access. To activate the move, call the stored procedure ADMIN_MOVE_TABLE.

Some of the other features available are; the maximum size of large and temporary tablespaces have been increased, there is also support for distribution map for larger data warehouses. Database partition servers will now be online as soon as they have been added to an instance and the DESCRIBE command will provide additional information for indexes.

XML

PureXML was first introduced within version 9.1 and has now been developed further to incorporate more flexibility and reliability. XML has been enhanced to further its capabilities and usage within partitioned tables, allowing users to run queries on XML data. XML columns can be added to multidimensional clustering tables, increasing performance for both indexes on XML data and for MDC indexes. Its uses can also be beneficial within declared global temporary tables, user-defined functions and utilities can be run against tables with XML columns such as runstats.

Maximising storage within the database is a target that is achieved using the compression capabilities of DB2 and the flexibility of XML data compression at row level, previous versions did not allow for this type of compression. With version 9.7 a compression dictionary will be created when the data in the XML storage object is compressed. The XML columns must be created at version 9.7 and row compression must be set for the table, in order to make this change, use the ALTER TABLE/CREATE TABLE commands and set to COMPRESS YES. If you attempt to compress a previous DB2 versions XML data, it will fail. A work around is to migrate the table by using the ADMIN_MOVE_TABLE stored procedure.

The ***Describe command***, can now be used for XML indexes providing detailed information. With the ***Alter table statement***, you can now specify a REORG statement for a table that contains XML columns. An increased level of concurrency can be achieved when running a reorg of indexes or during the creation of indexes on XML data, version 9.7 allows transactions to insert, update and delete data within a table. ***Predicate pushdown*** is available to both SQL/XML and XQUERY statements, this optimisation function allows data to be filtered earlier and promotes better index usage***. XML parsing*** creates diagnostic information when it is passed. A stored procedure XSR_GET_PARSING_DIAGNOSTICS provides the error handling, previously provided within version 9.5, this version provides more detailed messages***.***

High Availability

A key requirement by many organisations is 24x7 availability of their systems The internet has allowed many businesses to move online and create a further revenue stream. However, this requires their systems to be available around the clock. DB2 has become a front runner in high availability and each new version sees more enhancements to an already superior offering. DB2 version 9.7 again incorporates some more enhancements for high availability.

DB2 has become a front runner in high availability and each new version sees more enhancements to an already superior offering.

Enhancements:

- Clustered management software integration support now available for Solaris.
- DB2 is more stable during critical errors and is able to still function safely during these times, whilst still providing detailed diagnostic information. Additional messages have been added to the administration log for database administrators.
- The ability to specify the amount of disk space taken up by the administration notification and diagnostic logs using the ***diagsize*** database manager configuration parameter.

Online Schema Change

An increased availability of database objects for users in DB2 9.7 is provided through 'online schema changes'. Prior to this release much of the functionality did not allow a user the continued availability of a database object for both read and write, more often than not a user could read but not write, whilst a specific utility, alter command, create was running. The new changes now allow a user to change column names online using the ALTER table statement, there is more support for ALTER column set data type and tables can be moved or changed online using ADMIN_MOVE_TABLE. The new changes are immensely beneficial to DBAs and help to reduce any significant database downtime and allow the database environment and its objects to be available to users at all times.

Monitoring Enhancements

The new monitoring enhancements in 9.7 can be used to diagnose issues within your environments, using table functions and new event monitors. The new functions and event monitors are much improved on the existing monitors and snapshots.

Some of the benefits of the new monitoring enhancements cover;

New monitoring interfaces: these new interfaces are SQL based and are there for reporting and monitoring purposes, specifically the database environment, objects and the package cache. Reports can detail any issues or work for the environment, database objects, tablespaces and their containers as well as any SQL within the package cache. These new interfaces are far more efficient than existing monitoring tools and have less of a performance impact on the environment. The new table functions that have been incorporated fall under three levels; **system, activity** and **data object.**

Increased granular monitoring: Within version 9.7 as standard, the new monitor elements allow for more efficient monitoring through new database configuration parameters. SQL can be used to query the new elements without a huge performance overhead. An increased level of control is noted when gathering data, the information collected is specific to what you require and reported via table function

interfaces at three levels; **system, activity and data object level.** There are eight new parameters that have been added, where the defaults provide a minimal output.

Dynamic and Static SQL monitor: the new parameter MON_GET_PKG_CACHE_STMT is used to report on dynamic and static SQL queries, piping out information such as slow running queries, resources used, locks etc. Previously only snapshots for dynamic statements were available, now this new parameter allows analysis of both dynamic and static SQL.

UOW monitoring: unit of work monitoring is a new event monitor which has replaced the previous transactions event monitor. Designed to be more efficient than it's predecessor, it carries out analysis through the use of table functions **mon_get_unit_of_work and mon_get_unit_of_work_details.** The data can also be retrieved from the event monitor using XML documents table function, using **evmon_format_ue_to_xml** or by using a table through the new procedure **evmon_format_ue_to_tables.** Once the data has been captured, it is presented within the event table and will require formatting. The data can be converted into both XML or text documents, using db2evmonfmt.

Monitor time: better facilities to monitor the time spent by DB2 allow the DBA to analyse where DB2 is spending its time and using its resources, enabling the DBA to locate and resolve any issues. The time monitor will provide data on processing time, wait times, lock waits, buffer pool activity, time spent outside of DB2. The overall picture is time based and how much time certain aspects of DB2 take when processing. Importantly the new monitor also identifies if the slowdown is outside or inside DB2. This is certainly important as the database is always to blame for slowdowns and wait times; at least this monitoring will help to eliminate one or the other.

Locks: A whole host of event monitors that refer to locking and deadlocks have been removed (DB2_CAPTURE_LOCKTIMEOUT, DB2DETAILDEADLOCK, CREATE EVENT MONITOR FOR DEADLOCKS) and replaced with a more efficient event monitor (CREATE EVENT MONITOR FOR LOCKING), which covers a host of monitors, such as deadlocks, lock timeouts and lock waits. Again unformatted data can be more readable using db2evmonfmt and converting into XML or text documents. A further table function can be used to access the data via XML documents using EVMON_FORMAT_UE_TO_XML or via a table using EVMON_FORMAT_UE_TO_TABLES.

SQL Capture: an activity event monitor can now capture certain segments of an SQL statement and the access plans for the segment can be extracted using db2expln. To make use of this attribute, attach the following command onto a service class; COLLECT ACTIVITY DATA WITH DETAILS SECTION. The advantage of gathering specific segments, means there is less data returned and only the required data is returned.

Key Business benefits

The most obvious business benefit of moving to DB2 9.7 is lower costs. This covers many areas including:

- Lower storage costs
- Lower data management and analysis costs
- Lower database administration costs
- Lower development costs

In addition to lower costs this new version of DB2 brings to the market heightened security features which are key for all businesses in operation today. With More than 70% of data breaches coming from internal weaknesses it is vital that organisations put internal security at the top of their "to do" list.

With the new distinction between duties functions in DB2 9.7 organisations can maximise the security of their key intellectual property.

Another major benefit for the business is the enhanced availability which can be achieved with DB2 9.7. DB2 is widely recognised as a stable platform but the developments in this version have taken stability to a new level meaning that organisations can achieve 24/7 availability to their critical DB2 data.

Conclusion

With the release of DB2 9.7 we have seen IBM make an already good product into a genuinely great one. With it's many benefits, both to the business and technical teams, DB2 9.7 is the obvious choice for those looking to achieve lower costs whilst heightening security and reliability. IBM have pushed the boundaries and really taken the fight to Oracle and now have a compelling reason for organisations to change allegiance.

Triton Consulting are Data Management specialists and IBM Premier Business Partners. Specialising in DB2 for both the mainframe and distributed systems, Triton provide a full range of services from consultancy through to education and 24/7 DB2 support. Our Consultancy on Demand services allows organisations to purchase 10, 20 or 50 hours support which can be used for training, consultancy or to cover absences within the existing DBA team. For more information on Triton and the solutions they provide visit www.triton.co.uk