



Data Set Closing

From Deferred To Pseudo

By Craig S. Mullins

The CLOSE option can be specified on DB2 indexes and tablespaces. It indicates whether DB2 should close the underlying VSAM data set when it is not in use. The valid options are YES or NO.

However, the ramifications of the CLOSE options are not quite so clear-cut. The concept of what constitutes "in use" is completely different, depending on which DB2 version and release you are using. Likewise, the actual technique used to close the underlying data sets varies considerably — to the point of changing which CLOSE option is more efficient.

The Original DB2 CLOSE

The default CLOSE option is, and always has been, CLOSE YES. Prior to DB2 V2.3, this caused numerous headaches because CLOSE YES meant that the underlying VSAM data set was to be opened and closed every time it was used.

This was highly inefficient and caused most shops to enforce coding CLOSE NO on all DB2 objects.

The DB2 V2.3 CLOSE

DB2 V2.3 implemented a de-

ferred close strategy for DB2 page sets. When a tablespace or index defined using CLOSE YES is no longer being accessed, DB2 performs a "logical" close. The logical close places the page set on a "drain" queue and updates SYSIBM.SYSLGRNG, but does not physically close the data set. The physical close does not happen until the number of page sets on the "drain" queue is 99 percent of DSMAX. DSMAX is set in the DSNZPARMs and for DB2 V2.3 can be up to 10,000.

When the maximum is met, page sets are removed from the "drain" queue and physically closed, so the least recently used page sets are closed first. Physical data set closing minimizes the amount of virtual storage required since each opened data set requires 2.2KB of storage. As log ranges are recorded for "logically" closed tablespaces, the SYSIBM.SYSLGRNG table is updated and performance degradation tends to occur.

In addition to the SYSLGRNG updating, a small amount of overhead is encountered with CLOSE YES, due to the additional code that DB2 must execute to perform the logical close. This additional overhead should be minute. However, because of the SYSLGRNG updating, few shops use CLOSE YES for DB2 V2.3 tablespaces.

CLOSE NO under DB2 V2.3 causes DB2 to leave unused data sets open until a STOP DATABASE command is entered or the

system terminates. This results in improved performance since the open and close is avoided every time the data set is used. Opening a data set takes approximately 0.7 seconds of CPU. A single complex query may require that many data sets be opened. Keep in mind that each tablespace and index space involved with a query has a data set that needs to be opened at a CPU cost of 0.7 seconds!

The DB2 V3 CLOSE

As of Version 3, DB2 has changed the process for opening and closing the underlying VSAM data sets. The end result essentially eliminates the differences between CLOSE YES and CLOSE NO.

The DB2 V3 deferred close process is fundamentally equivalent for both CLOSE YES and CLOSE NO data sets. Only one difference remains: the order in which the underlying data sets are physically closed when the open data set threshold is reached. This threshold is still 99 percent of the value specified for DSMAX (or the MVS limit for the number of open data sets). If either threshold is met, DB2 will physically close some of the open data sets (3 percent of DSMAX). This is accomplished by closing data sets until the 3 percent threshold is met as follows.

- First, DB2 will physically close data sets coded as CLOSE YES.
- When no remaining CLOSE YES candidates remain, DB2

A new feature of data set closing under DB2 V3 is when the VSAM open-for-update timestamp is set. This change eliminates logging for read-only data sets.

Pseudo Close DSNZPARMs		
Parameter	Definition	Default Value
PCLOSEN	Number of checkpoints between page set modification.	5 checkpoints
PCLOSET	Amount of time since page set modification.	60 minutes

Table 1

will physically close data sets marked as CLOSE NO.

Another new feature of data set closing under DB2 V3 is when the VSAM open-for-update timestamp is set. When opening a data set, DB2 V3 will not modify this timestamp until data in the page set is updated. In all prior releases, it was updated when the data set was opened. This change effectively eliminates logging for read-only data sets.


A data set is considered logically closed when an application is deallocated from the data set. This occurs at commit or deallocation time, depending on the RELEASE parameter specified at BIND time (i.e., COMMIT or DEALLOCATE). When a data set is logically

closed, the data set in-use count is reduced by one. When the data set in-use count reaches zero, the data set is considered not to be in use. Furthermore, DB2 V3 will not update SYSLGRNG when the data set is logically closed or when the use count for a page set reaches zero.

Additionally, two DSNZPARMs control a new type of DB2 V3 close processing known as "pseudo" close (refer to Table 1). If an updated page set is not modified during the period defined by PCLOSEN and PCLOSET, the page set will be set to a read-only state. This is called a pseudo-close and when it occurs, SYSLGRNG is updated. When a tablespace data set is pseudo-closed, the

SYSLGRNG entry is closed and any updated pages are externalized to DASD.

Conclusion

As a general rule of thumb, consider using CLOSE YES for objects in DB2 V3. For performance-critical or hyper-sensitive (the CEO accesses them at whim) objects, CLOSE NO can be specified so they always will be the last objects physically closed. In short, little remaining difference exists between the CLOSE YES and CLOSE NO parameters. 

ABOUT THE AUTHOR

Craig S. Mullins is the author of "DB2 Developer's Guide," Second Edition (ISBN 0-672-30512), SAMS Computer Publishing, Carmel, IN, which contains DB2 V3 development techniques, tips and guidelines. He is a researcher in the education department at Platinum Technology, Inc., specializing in DB2, Sybase, client/server and object-oriented courseware. Platinum Technology, Inc., 1815 S. Meyers Rd., Oakbrook Terrace, IL 60181, (800) 442-6861.

REDUCE CPU BY UP TO 40%!

MEET YOUR BATCH WINDOW EVERY NIGHT!

SAVE 10% - 25% SYSTEM RESOURCES ON EVERY DB2 BATCH JOB!

ELIMINATE -904 AND -911 ABENDS FOREVER!

Is your DB2 batch workload growing while your batch window shrinks? Simplify your batch DB2 and reduce both elapsed times and CPU times by implementing **Database Attach**. For many batch jobs, you'll see elapsed time savings of up to 25%. For jobs that utilize our new Locking Feature, you can see up to 40% CPU savings and up to 60% elapsed time savings. The patented Enqueue Serialization Feature* will eliminate -904 and -911 abends. The SQL Monitor will give your applications staff a batch monitor designed with their needs in mind. Now you can meet the challenges of a growing DB2 environment. No source code changes are required, so implementation is a snap. Call us today for more information or a free 30 day trial.



210 Ridgefield Court
Asheville, NC 28806

704-665-0100 • 800-669-7076

*Patent No. 5,369,764

**D
A
T
A
B
A
S
E

A
T
T
A
C
H**