
Rule DB2-226: DWQT, VDWQT, or VPSEQT might be too small

Finding: CPExpert noticed that a buffer pool was used for sort activity, and the DWQT, VDWQT, or VPSEQT thresholds were relatively low. This finding means that the DWQT, VDWQT, or VPSEQT thresholds might be too small.

Impact: This finding can have a LOW IMPACT or MEDIUM on the performance of the DB2 subsystem.

Discussion: Buffer pool write operations are usually performed concurrently with user requests. Updated pages are queued by data set until they are written when:

- A DB2 *checkpoint* is taken
- The percentage of unavailable pages in a virtual buffer pool exceeds a preset limit called the Deferred Write Threshold (DWQT).
- The percentage of updated pages in a virtual buffer pool for a single data set exceeds a preset limit called the Vertical Deferred Write Threshold (VDWQT).

Up to 32 4KB or 4 32KB pages can be written in a single I/O operation.

Please refer to Rule DB2-225 for a more complete discussion of the DWQT and VDWQT thresholds.

The VPSEQT threshold is a percentage of the virtual buffer pool that might be occupied by sequentially accessed pages. These pages can be in any state: updated, in-use, or available. Hence, any page might or might not count toward exceeding any other buffer pool threshold. The default value for the VPSEQT is 80%.

While it is normally suggested that the DWQT and VDWQT thresholds values be relatively low to avoid "spikes" in page write activity, this advice does not apply to buffer pools heavily used by sort processing. Instead, the advice is reversed for these buffer pools, to allow sort processing to process sorted data in central storage.

DSNDB07 is the work area for DB2 sorting (and other related activity). Not only will the pages in DSNDB07 normally be referenced quickly, but the references typically are sequential. If the DWQT and VDWQT were set low, the pages would be written to DASD when these thresholds were reached.

However, the pages would be quickly reloaded. Consequently, the buffer pool that contains the DSNDB07 work files should have the DWQT, VDWQT, and VPSEQT thresholds set very high¹ to improve sort performance.

CPEXpert examines the QBSTWFR variable (the number of sort merge passes) in the DB2STATB statistics. When the value of this variable exceeds the **QBSTWFR** guidance variable in USOURCE(DB2GUIDE), CPEXpert concludes that the buffer pool is heavily used for sort activity. CPEXpert produces Rule DB2-226 when (1) buffer pools are heavily used by sort activity, (2) the DWQT threshold is less than 90, (3) the VDWQT threshold is less than 90, and (4) the VPSEQT threshold is less than 90.

The default value for the **QBSTWFR** guidance variable is 1, indicating that CPEXpert should conclude that the buffer pool is used heavily for sort processing when merge passes occurred. **It is unlikely that this value is correct for every buffer pool in your installation!** A low value was selected simply to alert you to this consideration and to allow you to specify a value that is more relevant to the buffer pools in your installation.

The following example illustrates the output from Rule DB2-226:

```
RULE DB2-226: DWQT, VDWQT, OR VPSEQT MIGHT BE TOO SMALL

Buffer Pool 7: CPEXpert noticed that this buffer pool is used for
sort activity. The DWQT threshold was specified as 50, the VDWQT
threshold was specified as 10, and the VPSEQT threshold was specified
as 100. While it is normally suggested that the DWQT and VDWQT values
be relatively low to avoid "spikes" in page write activity, this advice
does not apply to buffer pools heavily used by sort processing.
Instead, the advice is reversed for these buffer pools, to allow sort
processing to process sorted data in central storage. This situation
occurred for Buffer Pool 7 during the intervals shown below:
```

MEASUREMENT INTERVAL	BUFFERS ALLOCATED	TOTAL MERGE PASSES	THRESHOLDS REACHED DWQT	VDWQT
15:20-15:20, 15SEP1999	5,000	1,647	7	270
15:50-15:50, 15SEP1999	5,000	1,782	0	17
16:19-16:19, 15SEP1999	5,000	1,660	0	277
16:49-16:49, 15SEP1999	5,000	1,486	0	102
17:19-17:19, 15SEP1999	5,000	1,104	0	151

Suggestion: If Rule DB2-226 is produced regularly, you should consider the following alternatives:

¹ This comment is valid only if the buffer pool is used exclusively (or extensively) for sorting.

-
- Verify whether the buffer pool is used heavily (or exclusively) for sorting. If this is the case, you should set the DWQT, VDWQT, and VPSEQT thresholds to very high values. If the buffer pool is used exclusively for sorting, you should set these values to 90.
 - You should alter CPEXPERT's analysis by modifying the **QBSTWFR** guidance variable in USOURCE(DB2GUIDE), for any buffer pools that are not exclusively (or heavily) used for sort activity.

Please note that CPEXPERT allows guidance variables to be specified globally (that is, they apply to all buffer pools) or for a specific buffer pool. A reasonable approach would be to set the global value for the QBSTWFR guidance variable to some high value (e.g., 999999999), but set QBSTWFR to the default value of "1" for any buffer pool that is used primarily for sort activity.

Alternatively, you can simply "turn off" Rule DB2-226 after you have decided on appropriate values for the DWQT, VDWQT, and VPSEQT parameters for the buffer pool used primarily for sort activity.

- You can increase the DWQT, VDWQT, and VPSEQT thresholds for buffer pools that are exclusively (or heavily) used for sort activity, by using the ALTER BUFFERPOOL command. Changing a threshold in one virtual buffer pool or hiperpool has no effect on any other virtual buffer pool or hiperpool. You might wish to set the DWQT threshold to 90, the VDWQT threshold to 90, and the VPSEQT threshold to 95. This would indicate that 95% of the use of DSNDB07 is sequential and that updated pages should be written to DASD only when the pool is almost totally full.

This alternative is valid only if the buffer pool is sufficiently large that these thresholds are not reached frequently. If the thresholds are reached frequently, system performance could suffer. In this case, you should consider increasing the size of the buffer pool.

Additionally, note that the Sequential Prefetch (SPTH) is set at 90, which means that sequential prefetch would be disabled when the percent of unavailable buffers reached 90% of the buffer pool. This situation would be highly undesirable for most buffer pools, but is not a concern for buffer pools that are dedicated to sort processing.

Reference: DB2 for OS/390 Version 3: Installation Guide
Section 7.5.1 (Tuning Database Buffer Pools)
Section 7.5.4 (Controlling Sort Pool Size and Sort Processing)

DB2 for OS/390 Version 4: Installation Guide
Section 5.4.1 (Tuning Database Buffer Pools)
Section 5.4.4 (Controlling Sort Pool Size and Sort Processing)

DB2 for OS/390 Version 5: Installation Guide
Section 5.4.1 (Tuning Database Buffer Pools)
Section 5.4.4 (Controlling Sort Pool Size and Sort Processing)

DB2 for OS/390 Version 6: Installation Guide
Section 5.4.1 (Tuning Database Buffer Pools)
Section 5.4.4 (Controlling Sort Pool Size and Sort Processing)

DB2 UDB for OS/390 and z/OS, Version 7: Administration Guide
Section 5.4.1 (Tuning Database Buffer Pools)
Section 5.4.4 (Controlling Sort Pool Size and Sort Processing)

DB2 UDB for z/OS Version 8: Administration Guide
Chapter 26. Tuning DB2 buffer, EDM, RID, and sort pools
Tuning database buffer pools
Controlling Sort Pool Size and Sort Processing

DB2 UDB for z/OS Version 9: Performance Monitoring and Tuning Guide
Chapter 4. Tuning DB2 buffer, EDM, RID, and sort pools
Tuning database buffer pools
Controlling Sort Pool Size and Sort Processing

DB2 10 for z/OS: Managing Performance
Chapter 7. Configuring storage for performance
Buffer Pool Thresholds
Chapter 11. Using EXPLAIN to improve SQL performance

DB2 11 for z/OS: Managing Performance
Chapter 1. Managing performance in general
Buffer Pool Thresholds
Chapter 45. Investigating SQL performance by using EXPLAIN

DB2 11 for z/OS Buffer Pool Monitoring and Tuning
