
Rule DB2-512: Log reads satisfied from active log data set

Finding: A relatively large percent of log reads were satisfied from active log data set.

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

Discussion: DB2 logs changes made to data, and logs other significant events as they occur. DB2 initially writes the changes to the *log output buffer*, the changes are then written to the *active log*, and eventually the active log is off-loaded to the *archive logs*.

The output log buffers are formatted as VSAM control intervals (CIs). The log records containing changes are placed sequentially in the output log buffers, and each log record is identified by a continuously increasing relative byte address (RBA).

The log output buffers are written the active log, which are VSAM data sets. Multiple active logs are defined, and the active logs are used sequentially; as an active log is filled, the full active log is off-loaded to an archive log and log writing is switched to the next active log. The active logs are recycled (when the last defined active log is full, the first defined active log is then used).

A new archive log data set is created each time that the active log is off-loaded.

The size of the log output buffer, the size and number of active log data sets, and the number of archive log files have a significant effect on DB2 logging performance. This rule (Rule DB2-512) analyzes a potential problem when log reads were satisfied from active log data set..

When a rollback or recovery action is required and data is backed out, for optimal performance the data should be available in the output buffer or in the active log. It is preferable that the data be in the output buffer, but this is not always possible with an active DB2 environment. However, if a large percent of reads are satisfied from the active log, you should ensure that the output buffer is as large as possible.

CPExpert computes the percent of reads satisfied from the active log by the following algorithm:

$$\text{Percent reads from active log} = \frac{\text{Reads from active log (QBSTRACT)}}{\text{Total reads from log}}$$

CPEXpert compares the resultant percent with the **PCTRACT** guidance variable in USOURCE(DB2GUIDE). CPEXpert produces Rule DB2-512 when the percent of reads satisfied from the active log exceeds the value specified by the **PCTRACT** guidance variable.

The default value for the PCTRACT guidance variable is 25, indicating that CPEXpert should produce Rule DB2-512 when more than 25% of the log reads were satisfied from the active log.

The following example illustrates the output from Rule DB2-512:

RULE DB2-512: LOG READS WERE SATISFIED FROM ACTIVE LOG DATA SET			
<p>The DB2 Log Manager statistics revealed that more than 25% of the log reads were satisfied from the active log data set. It is preferable that the data be in the output buffer, but this is not always possible with an active DB2 environment. However, if a large percent of reads are satisfied from the active log, you should ensure that the output buffer is as large as possible. This finding occurred during the intervals shown below:</p>			
MEASUREMENT INTERVAL	TOTAL LOG READS	LOG READS FROM ACTIVE LOG DATA SET	PERCENT
21:32-21:47, 08SEP1998	56,393	56,391	100.0

Suggestion: As mentioned above, it is preferable that the data be in the output buffer when a log read occurs. While this is not always possible with an active DB2 environment, you should attempt to minimize the reads to the active log or to the archive log. Reads to the output buffer do not require an I/O operation, while I/O is required for reads to the active log or archive log.

If a large percent of reads are satisfied from the active log (and thus DB2 is incurring I/O operations), you should consider the following alternatives:

- Ensure that the output buffer is as large as possible. The larger the output buffer, the more likely a requested RBA can be found without a read request to either the active log or the archive log. The OUTPUT BUFFER field on panel DSNTIPL specifies the size of the output buffer used for writing active log data sets. The default size of OUTBUF is 40000KB, and the maximum size is 400000KB. The buffer must be large enough to prevent buffer shortages in order to decrease the number of forced I/O operations (forced because there are no more buffers) or wait conditions.

IBM recommends that you chose the maximum value that your system can tolerate for this field.

- Whenever you allocate new active log data sets, preformat the log data sets using the DSNJLOGF utility described in Section 3 of Utility Guide and Reference. This avoids the overhead of preformatting the log, which normally occurs at unpredictable times.
- Review and apply IBM's suggestions in the "Controlling the Log" subsection (in the "DB2 Logging" section referenced below) of the *DB2 Administration Guides*.
- You can design the DB2 log for better performance. In particular, try to:
 - Minimize device contention on the log data sets by placing data sets correctly.
 - Avoid waits that occur because no log buffer is available.
 - Define enough active log data sets to prevent DB2 from waiting while a log is archived.
 - Make the active log large enough that backouts do not have to use the archive log.
- Place the copy of the bootstrap data set and, if using dual active logging, the copy of the active log data sets, on volumes that are accessible on a path different from that of their primary counterparts. DB2 writes serially or in parallel to the log data sets. If a 4KB log control interval is written for the first time to DASD, the write I/Os to the log data sets are done in parallel. If the same 4KB log control interval is again written to DASD, then the write I/Os to the log data sets must be done serially to prevent any possibility of losing log data in case of I/O errors on both copies simultaneously. This improves system integrity; there is no I/O overlap in dual logging, except when the WRITE THRESHOLD value is reached, triggering full output log buffers to be written.
- You can alter CPEXPERT's analysis by modifying the PCTRACT guidance variable in USOURCE(DB2GUIDE).

Reference: DB2 for OS/390 Version 5: Administration Guide
Chapter 4-3. Managing the log and the bootstrap data set
Section 5.5.5 (DB2 Logging)

DB2 for OS/390 Version 6: Administration Guide

Chapter 4-3. Managing the log and the bootstrap data set

Section 5.5.5 (DB2 Logging)

DB2 UDB for OS/390 and z/OS, Version 7: Administration Guide

Chapter 4-3. Managing the log and the bootstrap data set

Section 5.5.5 (DB2 Logging)

DB2 UDB for z/OS, Version 8: Administration Guide

Chapter 17. Managing the log and the bootstrap data set

Chapter 27. Improving resource utilization (DB2 Logging)

DB2 for OS/390 Version 6: Administration Guide

Section 5.5.5 (DB2 Logging)

DB2 UDB for OS/390 and z/OS, Version 7: Administration Guide

Section 5.5.5 (DB2 Logging)

DB2 UDB for z/OS, Version 8: Administration Guide

Chapter 17. Managing the log and the bootstrap data set

Establishing the logging environment

Chapter 27. Improving resource utilization

DB2 Logging

DB2 UDB for z/OS, Version 9: Administration Guide

Chapter 16. Managing the log and the bootstrap data set

Establishing the logging environment

DB2 UDB for z/OS Version 9: Performance Monitoring and Tuning Guide

Chapter 5. Improving resource utilization

DB2 Logging

DB2 10 for z/OS: Administration Guide

Chapter 17. Managing the log and the bootstrap data set

Managing the log

DB2 10 for z/OS: Managing Performance

Improving log write performance
