Rule CIC406: VSAM data set might not be good candidate for shared data table

Finding: The CICS interval statistics showed that there was *very* high VSAM source data set activity for a CICS-maintained shared data table.

Impact: This finding has a MEDIUM IMPACT or HIGH IMPACT on the performance of the CICS region.

Logic flow: This is a basic finding, based on an analysis of the CICS interval statistics.

Discussion: An application can specify that a CICS VSAM key-sequenced data set (KSDS) file is to use shared data table services. When the file is opened, this specification causes CICS to copy the contents of the file into an *MVS data space*. The records can be accessed in an MVS data space significantly quicker than records read from the VSAM data set or via reads serviced by a Local Shared Resources (LSR) pool.

Since a major benefit of a data table is that records in the data table can be accessed quickly, this benefit is available only if the records actually are accessed in the data table rather than in the VSAM source data set. With a CICS-maintained data table, any operation requiring access to the source data set reduces the efficiency of the data table.

CICS supports two types of data table :

- **CICS-maintained data tables**. A CICS-maintained data table is one that CICS keeps in synchronization with their source data sets. That is, any update or delete action on a record in the data table is *automatically applied* to the source data set *before* being applied to the data table.

- **User-maintained data tables**. A user-maintained data table (UMT) is one that is not maintained by CICS, but is completely maintained by user code. A UMT is detached from its source data set after the table is loaded from the source data set, and changes made to the UMT are *not* reflected in the VSAM source data set.

For CICS-maintained data tables, the following file control commands access the VSAM source data set:

- READ commands with the UPDATE or RBA options.

- STARTBR, RESETBR, READNEXT, and READPREV commands with the RBA option.

- ENDBR command for a browse sequence that has accessed the source data set.

- READ and BROWSE commands (that would usually access only the data table) that find a gap in the key sequence of records in the data table. This gap might indicate that one or more records are missing from the data table. IBM documentation lists reasons that one or more records might be missing from the data table:

  - Records had been suppressed by the XDTRD global exit during table loading or by the XDTAD global exit when records were added to the table.

  - The maximum number of records (specified by the MAXNUMRECS value in the DEFINE FILE command) had been reached.

  - Insufficient virtual storage was available for the data table.

  - Some abnormal event occurred that prevented a record from being in the data table.

- READ, READNEXT, and READPREV commands for records that are currently being processed by a WRITE, REWRITE, or DELETE command. These commands need to first access the data table to determine that this situation exists.

- WRITE, REWRITE, and DELETE commands. These commands first update the source data set, and then attempt to update the data table.

- READ requests that encounter a "record not found" condition must access the VSAM source data set to retrieve the record not in the data table.

If a large percent of file control commands access the VSAM source data set, the benefits of using a data table can outweigh the overhead and virtual storage costs of maintaining a data table.

Shared data table statistics are available in MXG file CICFCR. CPExpert uses data in CICFCR to calculate the percent of file control commands that accessed the VSAM source data set[1], using the following algorithm:

$$\textit{Percent VSAM source data set accesses} = \frac{\textit{VSAM data set accesses}}{\textit{All file accesses}}$$

[1] Note that this algorithm would produce results only for CICS-maintained data tables. There would be no corresponding references to the VSAM source data set for user-maintained data tables.

where

VSAM data set accesses = A17DSGU+A17DSWRU+A17DSWRA+A17DSDEL + A17RMDEL + A17DSBRU - A17DTAVR

All file accesses = A17DSRD+A17DSGU+A17DSBR+A17DSWRU+ A17DSWRA+A17DSDEL+A17RMDEL+A17DSBRU- A17DTAVR

Please refer to the CICS Performance Guides for a description of the individual variables.

Note that removing the A17DTAVR value from the accesses accounts for the initial table loading process.

CPExpert produces Rule CIC406 when **more than 90%** of file control commands accessed the VSAM source data set, and this situation existed for **more than 75%** of the CICS statistics intervals being analyzed.

Before producing Rule CIC406, however, CPExpert applies the **MINSDTIO** variable to ensure that Rule CIC406 is produced only for shared data tables with a reasonably high amount to of I/O activity. The default value for the MINSDTIO is 500, indicating that Rule CIC406 will be suppressed unless at least 500 VSAM file control commands are issued to the shared data table.

Suggestion: For data tables to be effective, a significant percent of the file control commands must access only the data table rather than cause CICS to access the VSAM source data set. If Rule CIC406 is produced regularly, you consider the following alternatives:

- Determine whether records were not found in the shared data table, requiring access to the VSAM source data set. Rule CIC402 would normally be produced if records were not found in the shared data table. Records might not be found because number of records was limited by the MAXNUMRECS variable, because records were excluded by the XDTRD global exit while the table was loaded, or because records were excluded by the XDTAD global exit while records were dynamically added to the table. If records were not found, follow the advice given by Rule CIC402.

- You should consult with applications personnel to determine whether the VSAM file control commands can be modified to use less restrictive access techniques.

- You should consult with applications personnel to review the application

logic and coding. It is possible that an error exists with the application.

- If the high access to the source data set is a natural result of the application's correct logic, perhaps the file is not suited for a data table. This conclusion would become more certain if a *high* percent of file accesses required that CICS access the VSAM source data set. If more than 90% of the file accesses required that CICS access the VSAM source data set, the file probably should not be included in a data table.

- You can change the **MINSDTIO** guidance variable in USOURCE(CICGUIDE) if you want to have the CIC406 logic applied only for shared data tables with a larger number of I/O accesses.

Reference: CICS/TS for OS/390 Release 1.3 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

CICS/TS for z/OS Release 2.1 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

CICS/TS for z/OS Release 2.2 *CICS Shared Data Tables Guide*:
Section 6.1 Using the DEFINE FILE command to define data tables
Section 8.4 Interpreting data table statistics

