
Rule WLM059: Insufficient local page data sets are defined for migration

Finding: CPExpert has determined that the number of local page data sets is insufficient to allow parallel I/O operations for migration of pages from expanded storage to auxiliary storage.

Impact: This finding can have a LOW IMPACT, MEDIUM IMPACT, or HIGH IMPACT on performance of your computer system. The level of impact depends upon how often migration occurs.

This finding does not apply with z/OS running on z/Architecture (64-bit systems), since expanded storage is not used in this environment.

Logic flow: The following rule causes this rule to be invoked:
 Rule WLM400: Page-in from auxiliary storage was a major performance problem

Discussion: Please see the discussion contained in the referenced "Expanded Storage Management with MVS/ESA" regarding expanded storage operation. The discussion in this document is important for background information on the Migrator and other expanded storage concepts. Some of the relevant discussion is repeated below.

When the Migrator runs, it has a "migration quota" which it is trying to fill. This is simply the number of expanded storage frames which must be made available to alleviate the shortage. The quota of frames is the RCEAECOK threshold minus the RCEAECLO threshold. This quota differs depending upon whether the MVS version is SP4.2 or prior versions, and differs depending upon whether the Migrator was scheduled by the RSM or by the SRM.

- For pre-SP4.2, the initial low value (RCEAECLO) is 50 frames and the initial OK value (RCEAECOK) is 100 frames. For pre-SP4.2, the Migrator has a quota of at least 50 frames (100 - 50) which it will attempt to migrate.
- With SP4.2, the initial values are based on the MCCAECTH keyword in IEAOPTxx. The MCCAECTH has a default specification of (150,300). Thus, the default low value (RCEAECLO) is 150 frames and the OK value (RCEAECOK) is 300 frames. With SP4.2, the Migrator has a quota of at least 150 frames (300 - 150) which it will attempt to migrate.

-
- The initial RCEAECLO and RCEAECOK values can be adjusted based upon how often expanded storage becomes constrained. The RCEAECLO and RCEAECOK values can be increased when expanded storage is constrained (so the Migrator will run earlier). These values subsequently can be lowered if expanded storage is no longer constrained (so the Migrator does not run as soon). The RCEAECLO and RCEAECOK values will not be lowered below their initial values. The migration quota will never be less than the amount computed based on the default values.
 - The quota which the Migrator is attempting to fill can be larger than the default. This situation occurs if the SRM has scheduled the migration to accommodate the processor storage requirements of an address space which the SRM is attempting to swap in. The SRM can raise the RCEAECLO and RCEAECOK thresholds if necessary to make available more processor storage to contain the working set of a swapped-in address space.

It is important to appreciate that the Migrator does not migrate pages at a steady rate as might be reported by RMF, but migrates in "bursts" of pages. The "burst" will be the quota which the Migrator is attempting to fill. As the previous paragraphs discussed, the Migrator will have a quota of either 50 or 150 pages (or more), depending upon the version of MVS. This number of pages will be migrated during one "burst" of migration, even though RMF might report an average migration rate of only 10 pages per second (for example).

When the Migrator sends page-out requests to the Auxiliary Storage Manager (ASM), the Migrator sends the pages in page groups. The ASM will send these page groups out to local page data sets as parallel I/O operations, so long as there are sufficient local page data sets defined. As described above, the quota which the Migrator is attempting to reach is a minimum of 50 pages (pre-SP4.2) and 150 pages (with SP4.2), and this quota may be increased depending upon normal SRM adjustments.

A page group is 30 pages. Consequently, in order to take advantage of the significant performance benefits of parallel I/O operations for migration, a minimum of 2 local page data sets ($50/30=2$) should be defined for pre-SP4.2 environments. A minimum of 5 local page data sets ($150/30=5$) should be defined with SP4.2 and subsequent versions of MVS.

These are the **minimum** numbers based on the initial (default) values for the thresholds. However, the algorithms used by the SRM to adjust the RCEAECLO and RCEAECOK thresholds can increase the difference between the RCEAECOK and RCEAECLO values, depending upon how often expanded storage is constrained. Consequently, CPExpert believes

that at least 3 local page data sets should be defined for pre-SP4.2 environments and at least 6 local page data sets should be defined for SP4.2 and subsequent versions of MVS.

From a performance viewpoint, it is important to remember that the migrated pages will be sent in page groups, and these page groups place a significant load on the local page data sets. Significant page-in delays could occur if page-in operations should be queued behind migration page groups.

The page-in delays might not be reflected in "average" data (for example, the delays might not be revealed by average TSO trivial response times over some RMF measurement interval). However, individual page-in operations (and individual TSO responses) could be significantly delayed. These delays could become noticeable even in average data if sufficient migration occurs during any interval.

To minimize the periodic performance problems, sufficient local page data sets must be defined.

CPEXpert has detected that (1) a service class was significantly delayed because of paging from auxiliary storage, (2) fewer than the desired minimum local page data sets have been defined, and (3) more than a nominal amount of migration occurred.

Rule WLM059 is produced if less than six local page data sets have been defined.

It is possible that six local page data set is still too few local page data sets. However, SMF records contain no information which can be used to determine whether MVS has increased the migration quota above the default such that more than six local pages data sets would be optimal.

The following example illustrates the output from Rule WLM059:

RULE WLM059: INSUFFICIENT LOCAL PAGE DATA SETS ARE DEFINED FOR MIGRATION			
The number of local page data sets is insufficient to allow parallel I/O operations for migration of pages from expanded storage to auxiliary storage. If migration occurs often, defining too few local page data sets could have a significant impact on paging performance. During the measurement intervals shown below, migration was relatively high and at least one service class missed its performance goal because of paging delays from auxiliary storage. CPEXpert suggests increasing the number of local page data sets.			
MEASUREMENT INTERVAL	AVERAGE MIGRATION RATE	LOCAL PAGE DATA SETS ALLOCATED	SUGGESTED ALLOCATION
7:30- 8:00, 16AUG1995	1.8	3	6
8:00- 8:30, 16AUG1995	1.7	3	6

Suggestion: CPExpert suggests that you consider increasing the number of local page data sets so that migrated pages can be sent using parallel I/O operations. These local page data sets must be on separate paths, or the parallel I/O operations cannot be implemented.

It is not necessary that the local page data sets be the only data set on their volumes. Unless you have severe migration problems, it is unlikely that there will be a serious conflict between (1) the migration or subsequent page-in operations and (2) relatively low-activity data sets residing on the same volume as the local page data sets. It normally is far more important to allow the parallel I/O operations to be effective than it is to restrict the number of local page data sets simply because they must share the volumes with other data sets.

Reference: MVS/ESA SP4.2 Working Set Management and Block Paging Presentation Guide, GG66-3204, pages 50-54.

MVS (Real Storage Manager) source code, module IAREA.

"Expanded Storage Management with MVS/ESA", Deese, Donald R. (Computer Management Sciences), *CMG '93 Conference Proceedings*, The Computer Measurement Group, Inc., Chicago, IL.