
Rule WLM061: Swap data sets are defined

Finding: CPExpert has determined that swap data sets are defined in the swap/page configuration.

Impact: This finding can have a MEDIUM IMPACT on performance of your computer system. This rule applies only to an environment with expanded storage. This rule does not apply with OS/390 V2R10, as swap data sets are not supported beginning with this release of MVS.

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

Discussion: Swap data sets can be used by the Auxiliary Storage Manager (ASM) to contain Local System Queue Area (LSQA) and private area pages that are swapped in with the address space.

For systems with expanded storage, the RSM and SRM may divide the working set pages into a primary and secondary working set¹.

- **Primary working set.** The **primary working set** consists of LSQA pages, fixed pages, and one page from each virtual storage segment that is included in the working set².

The primary working set may be sent to expanded storage or may be migrated from expanded to auxiliary storage.

- The primary working set may be migrated to swap data sets if swap data sets are defined and if sufficient space exists on the swap data sets.
- If swap data sets are not defined or if insufficient space exists on the swap data sets, the primary working set is migrated to local page data sets.

¹This division is done only if the swap is to be done to expanded storage. If the swap is to be directly to auxiliary storage, the division is not done (a swap directly to auxiliary storage is called a **single stage swap**).

²The working set is composed of those address spaces with UIC of zero or one (and potentially an "enriched" working set with UIC greater than one if storage is not a constraint). A virtual storage segment is one megabyte of virtual storage.

-
- **Secondary working set.** The **secondary working set** consists all working set pages not included in the primary working set. These are most non-LSQA, non-fixed, working set pages. Notice that the secondary working set does not include swap trim pages³

The primary working set may be sent to expanded storage or may migrated from expanded to auxiliary storage. The secondary working set will always be migrated to local page data sets.

There are several advantages to using **only** local page data sets, rather than a mixture of swap data sets and local page data sets.

- The ASM load balancing algorithm selects the local page data set with the best performance to receive a page group. This algorithm automatically helps correct performance problems if local page packs are on heavily loaded paths or if local page packs are not dedicated. The ASM does not apply the load balancing algorithm to swap data sets.
- With expanded storage, most of the migration paging (that is, the migration of the secondary working set and migration of swap trim pages) is automatically sent to local page data sets. Thus, most of the pages associated with a swap (either directly in the case of the secondary working set, or indirectly in the case of swap trim pages) will be sent to local page data sets regardless of whether swap data sets are used. Consequently, swap data sets tend to be under-utilized in an expanded storage environment.
- Overall system performance normally would be much better if the volumes which were defined as swap data set volumes were redefined for local page data sets. The local page data sets would individually have a lower average page rate since there would be more volumes available (that is, the paging load would be spread over more volumes).

For example, suppose you had defined four local page data sets and two swap data sets. Performance would normally be significantly improved if you redefined the swap data sets as local page data sets, for a total of six local page data sets.

The following example illustrates the output from Rule WLM061:

³Swap trim pages are those pages trimmed from an address space before it is swapped out. The swap trim pages are the pages in central storage at swap time, which are not included in the working set. The swap trim pages may be sent to expanded if they meet the expanded storage criteria or they will be sent to auxiliary storage.

RULE WLM061: SWAP DATA SETS ARE DEFINED

CPEXpert has determined that swap data sets are used for swapping. Most performance analysts agree that swap data sets generally should not be used for swapping. Rather, local page data sets should be used in most environments. During the period being analyze, at least one service class missed its performance objective, and a major cause of delay to the service class was waiting for pages from auxiliary storage or a swap-in delay.

Suggestion: CPEXpert suggests that you consider changing your swap data sets to local paging data sets.

There may be circumstances in which you should retain the swap data sets. For example, you may have very large swap sets in an environment without adequate expanded storage. You may wish to retain swap data sets to prevent critical page-in operations from being slowed by the I/O required to service large swap sets.

The following issues should considered:

- Delay of critical page-in operations is unlikely to exist in an expanded storage environment. Since only the primary working set may be migrated to swap data sets (unless the swap is a single stage swap), little advantage is gained by having swap data sets. That is, the secondary working set will always migrate to local page data sets and the secondary working set is usually significantly larger than the primary working set. Since only the primary working set would be migrated, only the primary working set would be effected by having swap data sets.
- You normally should have sufficient local page data sets such that the ASM can initiate swap-out I/O operations in parallel to local page data sets. If the I/O operations are initiated in parallel, then the maximum delay to page-in operations normally would be only the time required to transfer a page group (30 pages for local page data sets).

The time to transfer a page group normally would be about 50-60 milliseconds for an IBM-3380 paging device (the possible seek operation, search operation, and data transfer), and these times would become significantly less if the DASD were cached or if IBM-3390 devices were used for paging. This periodic delay would be offset if the swap data sets were converted to local page data sets, since more local page data sets would result in a lower average page-in time.

- Under some circumstances, the migration rate may be high. If the migration rate is high, one implication is that there are few available

pages in expanded storage. (The only purpose of migrating pages is because there is an insufficient number of available expanded storage pages.)

If there are few available expanded storage pages, the SRM will direct swaps to auxiliary as **single-stage** swaps, and will not prepare a primary and secondary working set. In this situation, allocating swap data sets may prevent the single-stage swaps from overloading the local page data sets.

Of course, if many swaps are sent to auxiliary rather than to expanded, you have basic problems with your expanded storage environment.

Samson (see references) gives additional examples in which you may wish to retain swap data sets.

Unless you have a very unusual situation, CPExpert suggests that you eliminate the swap data sets and (if appropriate) use the DASD volumes as additional local page data sets.

Reference: MVS Initialization and Tuning Guide
MVS/ESA: Section 2.4 (Performance recommendations)
MVS/ESA: Section 2.6 (question regarding expanded storage)

OS/390 MVS: Section 2.4 (Performance recommendations)
OS/390 MVS: Section 2.6 (question regarding expanded storage)

OS/390 MVS (V2R10): Changes

"Paging Analysis in an Expanded Storage Environment," Bretvas, Thomas (IBM Corporation), *Conference Proceedings*, 1987 International Conference on Management and Performance Evaluation of Computer Systems (CMG '87). **This paper is highly recommended!**

MVS Performance Management, Samson, Stephen L., McGraw-Hill Publishing Company (pages 241-242)

MVS System Resources Manager module IRARMST2.