

Software Capacity Planning: Migrating to 64-bit Mode

OS/390 Version 2 Release 10 and z/OS introduce support of 64-bit real storage environments. OS/390 V2R10 has the ability to run in either 31-bit or 64-bit mode on a z900. Installations migrating to z/OS on a z900 will only be able to take advantage of the 64-bit real storage environment. All installations migrating to new releases of the operating system need to do software capacity planning to correctly understand the effects of 64-bit mode real storage management on system utilization. This flash outlines the steps needed to do software capacity planning for a 64-bit mode migration.

Why introduce 64-bit mode real storage?

For processors on the scale of the zSeries, there generally will be some degree of storage constraint for any production workload when real storage is limited to 2GB, (31-bit addressing mode limit). Such a constraint manifests itself in the movement of pages to/from expanded storage. Processor cycles are consumed to manage this movement, thus reducing the cycles available to perform productive work. The cost of managing this constraint is directly related to the demand page rate to/from expanded storage. The demand paging rate varies by system image due to workloads and amount and type of processor storage (i.e. real and expanded) configured for the image.

IBM periodically reviews customer configurations and performance information looking for potential constraints due to large systems effects. The 2GB central storage limit was identified as a potential inhibitor to large systems growth, especially in light of the growth rates of new and traditional workloads forecasted for the zSeries environment. The investment in the introduction of the 64-bit mode architecture was made to allow the infrastructure to be ready to meet the projected zSeries capacity requirements.

Overview of 64-bit mode

In 64-bit mode, addressable storage can reach 2 exabytes, with the z900 processor supporting configuration up to 64GB of real storage. This increase in real storage eliminates the need to move demand pages to/from expanded storage (in fact, expanded storage does not exist in 64-bit mode). Hiperspace and VIO pages are simulated in a 64-bit mode environment and these pages are still moved. Processor cycles dedicated to managing storage in 31-bit mode can now be used for productive work in 64-bit mode.

The move to 64-bit mode is intended to be transparent to applications. All expanded storage facilities (for example, hiperspaces) continue to function even in an all-central storage environment. Workloads using hiperpools may choose to change to dataspace to increase efficiency. This is fairly easy to change and the data thus far in our labs indicate positive results from this change. Additionally, there are 64-bit mode releases of a number of software products, (like performance monitors), which must be installed. Check with your software vendors for any software changes needed to run in a 64-bit mode environment.

Capacity Planning for the Migration to 64-bit Real Storage

New workloads or increased growth in current workloads will also be able to take advantage of the new 64-bit mode real storage environments. Capacity planning requires the application of a balanced systems approach which requires the capacity planner to acquire not just increased CPU capacity for workloads, but also additional real storage and I/O capacity. The 64-bit mode real storage capability now allows customers to plan real storage capacity requirements without having to worry about the division of the processor storage resource between central and expanded storage. Customers migrating to a 64-bit mode real storage environment may also want to review their LPAR configurations. If the introduction of new LPARs was required due to central storage constraints, or existing LPARs could not be combined because of central storage contention, the introduction of an all real storage environment may remove this restriction.

There will be a small, additional cost to process work in 64-bit mode, due to increased instruction paths in the OS/390 and z/OS "Real Storage Manager (RSM)". Simply switching from 31-bit to 64-bit mode will invoke this cost. However, any demand paging eliminated by making this change will probably offset some or all of this cost. Where demand paging rates are significant in 31-bit mode, there will probably be a net benefit in capacity when moving to 64-bit mode. When demand paging rates are insignificant in 31-bit mode, there may be a small impact. In our lab workloads, we have seen up to a 20% improvement for highly constrained environments and up to a 5% loss for unconstrained environments. However, depending on your workload characteristics and amount of constraint, you will see various results.

Installations are encouraged to do capacity planning for both hardware and software migrations. As with most new software versions, the use of 64-bit mode needs to be evaluated in context with the business objectives for the zSeries platform.

Examples of workloads that will see performance improvements when moving to 64 bit are workloads with either high demand paging to/from expanded storage or workloads operating at low values of average high unreferenced interval counts (UIC). These workloads may see the CPU savings from eliminating paging and/or reduced storage management and the savings may offset the increased RSM costs incurred in 64-bit mode. Environments with these types of workloads should seriously consider moving to 64 bit mode.

Another environment where large benefits may be seen is a DB2 system where the transaction rate is limited by virtual storage constraints in the DB services address space, (max 2GB). In this case, some virtual storage savings may be realized by moving virtual buffers to data spaces, (which work very efficiently in an all-central storage configuration), enabling a higher DB2 transaction rate for the application. Again, environments with this type of workload should seriously consider moving to 64 bit mode.

In addition, when moving to z/OS which requires 64-bit mode, there are many new functional enhancements, (such as the Intelligent Resource Director and License Manager) which may be viewed as beneficial to the installation to outweigh any 64-bit real storage costs.

64-bit Mode Migration Tool Support

IBM's capacity planning processes are being updated to reflect our experiences with migration to 64-bit mode real storage. Information on the software migrations will be available through WWQ&A item RTA000017036, and via a web downloadable tool called SOFTCAP. This tool is available for download by IBM account teams at the following URL:

<http://w3-1.ibm.com/support/americas/wsc/cpsproducts.html>. It is also available to IBM business partners at the following URL: <http://www-1.ibm.com/partnerworld/sales/systems>. A customer version of the tool is under development and is planned to be available via download at URL <http://www.ibm.com/support/techdocs> in second quarter 2001.

The SOFTCAP tool will need as input your installation's storage contention information and storage management intensity level. The tool will integrate both metrics and will help installations gain the necessary software capacity planning information for doing a successful migration to 64-bit mode real storage. Information is provided below to help you to understand your installation's systems storage contention and storage management intensity level.

Understanding Your Environment

Customers will need to have an understanding of the paging environment on the OS/390 or z/OS images planning to be migrated to 64-bit mode. Capacity projections will be based on two system indicators; **storage contention** and **storage management intensity**.

Storage Contention:

Storage contention indicators are available from RMF (or equivalent). In this situation we will be using the demand paging rates, and the system's average high UIC count as indicators of storage contention. We expect that anyone with ≥ 5000 demand pages to and from expanded storage per second or with an average high UIC < 10 will need to seriously consider moving to 64 bit mode.

RMF paging reports can be generated or the RMF SMF record 71 can be post-processed to obtain the information on storage contention. For the RMF post-processor, reports can be generated by using RMF control cards:

```
REPORTS (PAGING)
```

Review the RMF reports or other performance reports and using the paging information calculate the demand paging rate:

$$\text{Demand Paging Rate} = \frac{\text{Total Pages} - (\text{Hiperspace Pages} + \text{VIO Pages} + \text{Shared Pages})}{\text{Time}}$$

The demand paging rate needs to be summed for movement to and from expanded storage. Below is an example of the type of RMF report you can use. The RMF report is broken into two segments to allow easier viewing. In this example the storage contention rate is calculated as:

Written To	Read From	Total
Demand Paging Rate	Demand Paging Rate	Demand Paging Rate
((144- (18 + 0 + 0)) +	(97- (0 + 0 + 0)))	= 223

These values indicates minimal storage contention. The determination of minimal storage contention is backed up by the average high UIC values.

```

OPT = IEAOPT00      MODE = ESA                      EXPANDED STORAGE MOVEMENT RATES
-----
HIGH UIC  MIGR AGE
MIN    254    85899
MAX    254    91168
AVG    254.0  88532.9

ESF CONFIGURATION
INSTALLED      ONLINE
-----
479232         479232

                WRITTEN TO      READ FROM      MIGRATED      FREED
                EXP STOR      EXP STOR      FROM          WITHOUT
TOTAL          RT          144.41      96.63      23.97      8.67
PAGES          %          100.0%     100.0%     100.0%
HIPERSPACE    RT          17.82      0.01      0.00
PAGES         %          12.3%     0.0%     0.0%
VIO           RT          0.13      0.19      0.00
PAGES         %          0.1%     0.2%     0.0%
SHARED        RT          0.00      0.00
PAGES         %          0.0%     0.0%

```

Storage Management Intensity:

Storage management intensity is an indicator of how much real storage management, (RSM), activity is being done by a workload. Such activity includes operating system services such as *getmain*, *freemain*, *page fault*, *page fix*, and *page free*. This activity is not controlled by an installation, rather it is a by-product of the type of workload the installation runs. Storage management intensity is not available as an external RMF metric, hence customers cannot directly measure the activity.

Since it is difficult to determine your actual storage management intensity, the input the SOFTCAP tool will use for your storage management intensity is levels of intensity, rather than exact numbers. SOFTCAP will use three levels of intensity; high, medium and low. Based on our current experiences we expect most workloads to fall in the low to medium levels. Using SOFTCAP, you will be able to choose the different levels (high, medium or low) and determine the effect on capacity. As additional experienced is gained with customer migrations to 64-bit mode, we may refine the levels, or give more guidance on which level to choose.

Special Notices

This publication is intended to help the customer manage a migration to 64-bit mode in either OS/390 V2R10 or z/OS environment. The information in this publication is not intended as the specification of any programming interfaces provided by OS/390 or z/OS. See the publication section of the IBM programming announcement for the appropriate OS/390 or z/OS release for more information about what publications are considered to be product documentation. Where possible it is recommended to follow-up with product related publications to understand the specific impact of the information documented in this publication.

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis without any warranty either expressed or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Performance data contained in this document was determined in a controlled environment; therefore the results which may be obtained in other operating environments may vary significantly. No commitment as to your ability to obtain comparable results is any way intended or made by this release of information.