

Large Memory Pages - Part 1

Fabio Massimo Ottaviani – EPV Technologies

April 2013

1 Introduction

Ever since the days when z/OS was called MVS, virtual memory has always been managed in 4096 byte pages. However, with the advent of z/OS 1.9 and z10 machines, 1MB pages can be used, whilst 2GB pages are also supported with z/OS 1.13 and zEC12.

The reason for this fundamental breakthrough is the exploitation of 64-bit architecture; it is now possible to create huge z/OS address spaces: up to 16 ExaBytes of virtual memory. In order to back this virtual memory, the more recent IBM hardware is able to provide up to 3 TB of real memory which is, by the way, much cheaper than previously .

In 31-bit mode the maximum address space size was 2GB; it could be mapped by using $256 * 2.048$ (524.288) 4K pages.

In 64-bit mode to map all the address space virtual memory $256 * 2.048 * 2.048 * 2.048 * 2.048$ (4.503.599.627.370.500) 4K pages would be required.

It's intuitive that managing such big address spaces with so many small 4K pages would not be very efficient; so to improve performance and to reduce CPU consumptions of memory-intensive workloads (e.g. DB2 and WebSphere applications), it is possible and advisable to use large memory pages.

In this paper after an introduction to virtual to real address translation, we will discuss what you have to do in order to exploit large memory pages and which metrics are available to analyse their utilization.