EPV Technologies Newsletter

30 June 2010   Number 06-2010

Editorial staff: Dino Gigli, Danilo Gipponi, Fabio Massimo Ottaviani (EPV Technologies) – Jon Olley (Inspired Solutions)

In this number

1) Tech Papers – Practical Capacity Planning for Memory – Part 1

2) Tech Support –

PM12256: DRDA Performance Improvement Using TCP/IP

OA26832: NEW FUNCTION FOR PROVIDING EXPANDED SERVICE UNIT VALUES IN THE SMF TYPE 30 RECORD

Past numbers of this newsletter are available on the web at http://www.epvtech.com

1) Tech Papers	**Practical Capacity Planning for Memory – Part 1**

The cost of mainframe memory has been substantially reduced during recent years; however its current cost is still in the order of some thousands of Euro per GB.

On the other hand the number of GBs supported by the newest z/OS releases has greatly increased making it possible to almost eliminate paging activity and to fully exploit the many available DIM (Data In Memory) techniques.

While it's true that z/OS can sustain high paging activity, paging and especially page faults are very disruptive to online applications performance. If a CICS region is getting 10 page faults per second this means that 10 transactions per second are stopped (or a transaction is stopped 10 times) and have to wait for the required pages to be loaded from page datasets. Furthermore you have to consider that CPU cycles are needed to perform paging. This is one of the sources of uncaptured CPU time.

Maintaining data in memory has two additional advantages: it improves performance by avoiding the I/O operations delay and it eliminates the CPU consumption still needed to perform I/O operations.

It's also important to remember that I/O operations can only be performed by standard CPUs and not by specialty processors (zAAP and zIIP) so the I/O reduction has positive effects on the software bill too.

All these factors mean that memory capacity planning is still a very important activity to perform.

The good news is that from a conceptual point of view the process is always the same:
- evaluate current capacity;
- estimate the utilization baseline;
- estimate the growth;
- forecast future resource needs.

The bad news is that dealing with memory, because of its nature, is more complicated than with other resources.

In this paper we'll discuss a practical methodology to do that.

If you want to receive the paper you can reply to this e-mail writing **"Practical Capacity Planning for Memory – Part 1"** *in the subject*

2) Tech Support

PM12256: DRDA Performance Improvement Using TCP/IP

Although this isn't marked as New Function, it provides a new method for DB2 to control the portion of SQL request workloads received via DRDA over TCP/IP connections. This allows more work (up to 60%) to be diverted to run on zIIP processors. It also reduces the overhead incurred for switching.

More details at: http://www-1.ibm.com/support/docview.wss?uid=swg1PM12256

OA26832: NEW FUNCTION FOR PROVIDING EXPANDED SERVICE UNIT VALUES IN THE SMF TYPE 30 RECORD

This APAR will provide expanded (to 8 bytes) fields equivalent to the existing 4 byte Service Unit fields in the SMF type 30 record. The existing fields are: SMF30SRV, SMF30CSU, SMF30SRB, SMF30IO, and SMF30MSO.
For long running jobs, the values in these fields may exceed 'FFFFFFFF'x, in which case the values will "wrap" to zero and continue to grow. Providing expanded, 8 byte, equivalent fields will allow for the values to grow past 'FFFFFFFF'x.

More details at: http://www-01.ibm.com/support/docview.wss?uid=isg1OA26832