EPV TECHNOLOGIES NEWSLETTER

May 2019

THIS MONTH HIGHLIGHTS

Measuring the CPU time:
misunderstood and unknown aspects
Part 2

IBM Software Pricing Announcement: "Tailored Fit Pricing for IBM Z"

Important IBM corrections

FORMULA OF THE MONTH

zIIP to CPU normalization in SMF

CPUs run at reduced speed in many IBM mainframe models while zIIP engines always run at full speed in any model.

There are a certain number of SMF metrics which provides the total CP time including both CPU and zIIP time. In these metrics the zIIP time is always normalized to the CPU speed.
So, if a certain activity used 1 CPU second and 1 zIIP second and the zIIP speed is 2 times the CPU speed, the cumulative metric value will be 3 seconds.

The general formula used is:
zIIP normalized-to-CPU time = zIIP time * (zIIP speed / CPU speed)

If you want to de-normalize the zIIP time to go back to the original time spent on the zIIP you have to invert the formula:

zIIP time = zIIP normalized-to-CPU time / (zIIP speed / CPU speed)

SMF 30 and 72 records provide coefficients in the SMF30SNF and R723NFFS that are a multiple of 256
They can be used to calculate the zIIP speed / CPU speed ratio as: SMF30SNF / 256 or R723NFFS / 256

Measuring the CPU time: misunderstood and unknown aspects – Part 2

CPU time is the most relevant performance metric measured at every z/OS site.
The reason why measuring the CPU time is considered so important is because hardware and software costs are generally correlated to the CPU usage or to the number of CPUs.

Nonetheless some aspects of CPU related metrics are often misunderstood or unknown.

In this paper we will discuss:
• CPU service unit's instability,
• CPU/zIIP time split and normalization issues in SMF 72,
• not accounted CPU time in SMF interval records,
• not accounted CPU/zIIP time in CICS and IMS transaction records,
• CPU/zIIP time split and normalization in subsystems measurements,
• CPU/zIIP consumptions for cryptography,
• how to avoid zIIP eligible time on CPU.

"**Measuring the CPU time: misunderstood and unknown aspects – Part 2**" in the subject

IBM Software Pricing Announcement: "Tailored Fit Pricing for IBM Z"

Tailored Fit Pricing for IBM Z is a transformational pricing option for IBM Z software. It offers simple, transparent, and predictable pricing for IBM Z software running on the z/OS(R) platform within a given country.

Tailored Fit Pricing introduces two comprehensive alternatives to the Rolling 4 Hour Average (R4HA)-based pricing model, for both new and existing workloads:

• The Enterprise Consumption Solution is a tailored consumption-based licensing model.
• The Enterprise Capacity Solution is a tailored full-capacity licensing model.

Both models dramatically simplify the existing pricing landscape, delivering flexible deployment options that are tailored to reflect the client's individual environments.

Both models include additional capacity for development and test environments as well as reduced pricing for all types of workload growth.

Complete announcement at:
https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS219-014

Important IBM corrections

OA56762: NEGATIVE VALUES CAN BE SEEN FOR SMF30_TIME_ZIIP_ON_CP IN TYPE30 SUBTYPE 3 RECORDS NEGATIVE VALUES CAN BE SEEN FOR SMF30_TIME_zIIP_ON_CP IN TYPE30 SUBTYPE 3 RECORDS ON Z/OS 2.2 WITH PTF UA97534 APPLIED.

The values for the job zIIP and zIIP_on_CP time are swapped when passed to SMF to cut the SMF30 record. The problem only exists at z/OS 2.2 with UA97534 applied.
More details at:
http://www-01.ibm.com/support/docview.wss?crawler=1&uid=isg1OA56762

OA56357: SP255 STORAGE ACCUMULATION FRAGMENTATION PPD SPD CDX INIT CPOOL
Users of CPOOL macro who define their cell pool to use a subpool whose characteristics have ownership of Job Step can result in the PPD / SPD not being freemained during task termination.
At z/OS 2.3, Contents Supervisor implemented a cell pool with CPOOL macro with these characteristics and it results in a buildup of PPD/SPDs for CDX pools.
The PPD/SPDs are found in subpool 255 and were building up notably in initiator address spaces. The CDX pools were using subpool 254 as their cell pool SP= parameter.
All job step owned subpools being used are exposed to this issue during task termination. (See Diagnosis Reference, Chapter 8: Storage Summary for other subpools that are owned by Job Step).
You may experience increase in cpu for batch jobs if there is significant buildup of control blocks.
More details at:
http://www-01.ibm.com/support/docview.wss?crawler=1&uid=isg1OA56357

RECIPES

Pasta con la crema di fave fresche

Ingredients

Fresh fava beans 1 kg
Pasta 400 g
Grated Pecorino 40 g
Almonds 30 g
thyme
lemon
Extra virgin olive oil
salt

Method

Make a puree with the shelled broad beans, pecorino cheese, almonds and 2 tablespoons of extra virgin olive oil. Let them cool quickly in cold water and peel them.

Puree the shelled broad fava beans with pecorino cheese, almonds and 2 tablespoons of extra virgin olive oil. Season with salt if necessary. Cook the pasta in boiling salted water and season with the cream of fava beans and a little cooking water from the pasta, if necessary.

Add thyme leaves and lemon

QUOTES
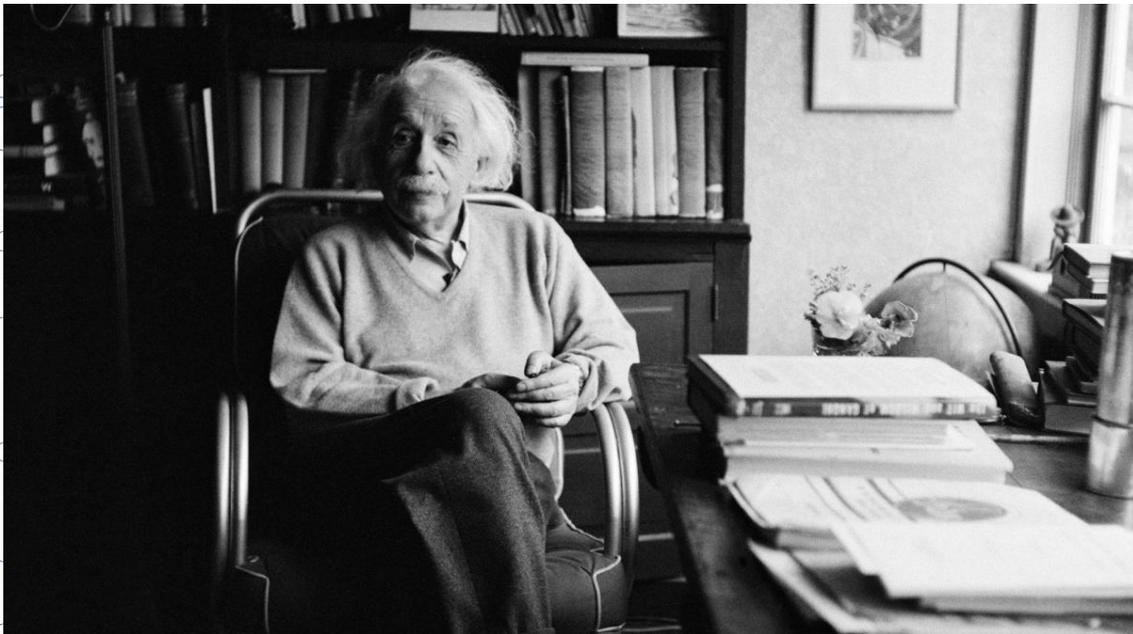
> *"Insanity is doing the same thing
> over and over again and expecting different results"*
> **Albert Einstein**

Our mailing address is:
EPV Technologies
Via Lugi Mancinelli, 106
Roma, RM 00199
Italy

Add us to your address book

Want to change how you receive these emails?
You can update your preferences or unsubscribe from this list.