



epv

IT Cost
Under Control

EPV Technologies

Newsletter

June 2023

THIS MONTH HIGHLIGHTS

- WLM controls of zIIP resources
- EPV User Group 2023

WLM controls of zIIP resources

Many years ago, IBM introduced the zIIP (z Integrated Information Processor) with the goal to provide capacity for “new” applications without increasing software costs.

Initially zIIPs were introduced to serve DDF applications but they rapidly

evolved to support Java applications (formerly served by zAAP, now obsolete), most of the functions provided by Db2 system address spaces (MSTR and DBM1), IPsec and many other products and functions, also provided by ISVs.

The last important step in zIIP evolution has been the introduction of z/OS Container Extensions (zCX) in z/OS 2.4 which enables clients to deploy Linux applications as Docker containers on z/OS as part of a z/OS workload. You must plan for the number of virtual CPUs to allocate to each zCX instance. zCX virtual processors can be dispatched on zIIPs or general purpose CPUs. If sufficient zIIP processors are available (greater than or equal to the number of zCX virtual processors), the majority of zCX instance processing can execute on zIIPs.

As you can imagine, WLM had to follow the described evolution to be able to control and manage the zIIP resources and the workloads using them.

In this paper we will discuss some of the important WLM concepts that every performance analyst should be aware of when looking at zIIP utilization.

If you want to receive the paper you can reply to this e-mail writing **"WLM controls of zIIP resources"** in the subject

EPV User Group 2023

The XXI EPV User Group will be in presence again. It will be held at Hotel Cicerone in Rome from 18th to 19th October,
The agenda structure will be:

Agenda Structure

11/10/2023

09:15 – 09:30	Introduction
09:30 – 10:00	EPV SOD
10:00 – 10:30	User presentation
10:30 – 11:00	User presentation
11:00 – 11:30	coffee break
11:30 – 12:00	EPV presentation
12:00 – 12:30	User presentation
12:30 – 13:30	Lunch
13:30 – 14:00	EPV presentation
14:00 – 14:30	EPV presentation
14:30 – 15:00	z/OS update
15:00 – 15:30	coffee break
15:30 – 19:00	Special Event
20:00	EPV user Group Dinner

12/10/2023

09:15 – 09:30	Introduction
09:30 – 10:00	EPV presentation
10:00 – 10:30	User presentation
10:30 – 11:00	User presentation
11:00 – 11:30	coffee break
11:30 – 12:00	User presentation
12:00 – 12:30	EPV presentation
12:30 – 13:00	EPV presentation
13:00 – 14:00	End of Conference & Lunch

The EPV User Group is a "not to miss" event for all Performance Analysts; it will give you the opportunity to share ideas with qualified experts and to listen to some of the EPV customers experiences. The most interesting features provided by the latest versions of all EPV products will also be presented.

The EPV User Group is free of charge and reserved to EPV customers. If you are not a customer yet but you are interested in participating, please answer to this e-mail asking for an invitation.

Mark these dates in your agenda to avoid missing this event.

More details in next newsletters.



Customer Questions

We need to understand who used the workfiles of a Db2 subsystem in a specific hour. In record SMF 101, I found these fields: QWAC_WORKFILE_MAX and QWAC_WORKFILE_CURR.

Can it make sense to use them as a starting point in the analysis?

EPV Technical Support answer

I would say yes. Also because in Db2 accounting data there is not much better.

Keep in mind that:

- QWAC_WORKFILE_MAX is the maximum number of 1K blocks used by the plan during the lifetime of the thread; therefore, if it started at 8 and finished at 10, the value can refer to any moment of the interval;
- QWAC_WORKFILE_CURR is the current number of 1K blocks used by the plan at thread close when the record is written.

In general, maybe it is better trying to use MAX.



BATCH_RCLMIGDS

The BATCH_RCLMIGDS parameter is included in the ALLOCxx member of the system parmlib.

It specifies the way migrated data sets are recalled.


Possible values are SERIAL and PARALLEL. SERIAL is the default.

With SERIAL, if a job step needs to recall three datasets, they will be recalled one at a time.

If you specify PARALLEL, they will be recalled in parallel. Using this option might reduce job execution time.

Our advice is to set it to PARALLEL.

Quotes



"...and there is nothing more beautiful than the moment before the trip, the moment when tomorrow's horizon comes to visit us and tell us about its promises"

Milan Kundera

Copyright © 2023 EPV Technologies, All rights reserved.

If you've received this mail by mistake, or you don't want to receive any more such messages, please send an e-mail to epv.info@epvtech.com with subject "REMOVE". You'll be promptly removed from the list. If you want to subscribe to this list you can do that simply by sending an e-mail to epv.info@epvtech.com with a subject "SUBSCRIBE".

If you've received this mail by mistake, or you don't want to receive any more such messages, please send an e-mail to epv.info@epvtech.com with subject "REMOVE". You'll be promptly removed from the list. If you want to subscribe to this list you can do that simply by sending an e-mail to epv.info@epvtech.com with a subject "SUBSCRIBE".

Our mailing address is:

EPV Technologies
Viale Angelico, 54

Roma, RM 00195
Italy

[Add us to your address book](#)

Our mailing address is:
EPV Technologies
Viale Angelico, 54
Roma, RM 00195
Italy

Images designed by : [Freepik](#), [Flaticon](#)