



**epv**

IT Cost  
Under Control

# EPV Technologies

## Newsletter

January 2022

### THIS MONTH HIGHLIGHTS

- [Measuring zCX applications - Part 2](#)
- [EPV Performance University 2022](#)
- [Installing and Configuring IBM Db2 AI for z/OS v1.4.0](#)

### Measuring zCX applications - Part 2

With z/OS 2.4, IBM introduced z/OS Container Extensions (zCX): a new z/OS 2.4 feature that enables customers to deploy Linux applications, packaged into Docker containers, which can run on z/OS with no change to the application code.

zCX instances are provisioned and managed through the z/OS Management Facility (z/OSMF).

Once provisioned, Docker users and application developers can operate within a zCX instance as they would in any other Linux-based, Docker environment.

In this paper, after a brief overview of zCX, we will describe the architecture of the EPV products when running on zCX.

The EPV “application” will then be used as an example in the next chapters to discuss:

- what are the useful z/OS commands you can use to control zCX instances;
- what you can do to measure zCX applications from “outside” (in z/OS) and from “inside” (in Docker/Linux).

*If you want to receive the paper you can reply to this e-mail writing **"Measuring zCX applications - Part 2"** in the subject*

---

## EPV Performance University 2022

---

Only one month to the 2022 edition of the EPV Performance University.

This is the current plan.

Course	Dates	Teacher
z/OS Performance Analysis	21-23 February 2022	Danilo Gipponi
MQ Performance Analysis	24 February 2022	Fabio Massimo Ottaviani
Db2 Performance Analysis	28 February - 1 March 2022	Fabio Massimo Ottaviani
WLM update	2-3 March 2022	Fabio Massimo Ottaviani

EPV products will be used as a map to make the path easier but most of the concepts discussed will be of general interest also for not EPV customers. Participation is free of charge for EPV customers, partners and invited guests. For all the others a subscription fee will apply.

All details and subscription forms are available at:  
[www.epvtech.com](http://www.epvtech.com)

---

## Installing and Configuring IBM Db2 AI for z/OS v1.4.0

---

“Artificial intelligence (AI) enables computers and machines to mimic the perception, learning, problem-solving, and decision-making capabilities of the human mind. AI development is made possible by the availability of large amounts of data and the corresponding development and wide availability of computer systems that can process all that data faster and more accurately

than humans can.

What happens if you infuse AI with a world-class database management system, such as IBM Db2? IBM has done just that with Db2 AI for z/OS (Db2ZAI).”

Redbook available at:

[Installing and Configuring IBM Db2 AI for IBM z/OS v1.4.0](#)



After an incident we used EPV to analyze HVCOMMON usage and concluded that the reporting is incorrect.

The incident (in short): the weekly stop/start of MQ and DB2 subsystems resulted in some DB2s failing to start due to HVCOMMON fragmentation. It failed to allocate a MEMOBJECT in COMMON 64-bit storage because there was no contiguous area of 5600MB available.

According to EPV reporting, we are using less than 2% of HVCOMMON, but that looks not correct.

I extracted the HVCOMMON related fields from the zPARSER DB and I used the formula given in the EPV for z/OS help.

It indeed matches the EPV reported values of about 1.5% usage.

By looking at yesterday data, I see that the SMF71CRA field, providing the number of frames backed in CS, shows that we used about 1GB of HVCOMM (with 66GB defined) but in another tool we get a different figure: it reports about 44GB of HVCOMMON in use.

Can you please review the HVCOMMON usage reporting?

***EPV Technical Support answer***

Based on the SMF data of the critical day you provided us, we can say that, as you told, the values in the SMF71C1A, SMF71CRA, SMF71CSA, SMF71C4A are correct. The percentage of HVCOMMON utilization reported in EPV is also correct.

We checked the RMF PAGING report produced with your and other customer reports and we saw that RMF provides exactly the same values as EPV.

However, after looking at your SMF data, manuals and research on the web about HVCOMMON, we found that there is an additional SMF field (SMF71CPA) which provides the average number of pages in-use (min and max are also provided in SMF71CPM and SMF71CPX). The strange thing is that values are much higher than the number of pages backed in Real, AUX and SCM.

Information about these fields is very poor. However, it seems that they also take into account fragmentation.

In general, it may be possible to have a small amount of HVCOMMON backed in REAL, AUX or SCM (and therefore a low HVCOMMON utilization) but a lot of fragmentation.

When Db2 (or eventually another subsystem) asks for a large chunk of contiguous HVCOMMON, it may happen that it cannot find it. In this case Db2 will not start.

This is what seems to be happened in your case when 5,6 GB of contiguous HVCOMMON were needed but they could not be found.

We saw, indeed, that less than 4,7 GB HVCOMMON were not in-use on that day.

By using this field, the HVCOMMON utilization, including fragmentation, could be calculated with the following formula:

$SMF71CPA / SMF71C1A$

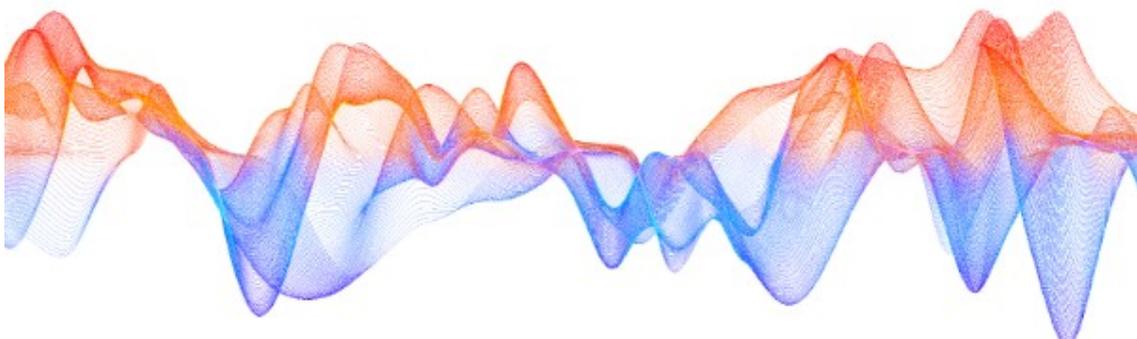
It is interesting to note that, if you run the DISPLAY VIRTSTOR,HVCOMMON you will get an answer like that:

```
D VIRTSTOR,HVCOMMON
IAR019I 18.25.46 DISPLAY VIRTSTOR 197
SOURCE = DEFAULT
TOTAL 64-BIT COMMON = 66G
64-BIT COMMON RANGE = 1982G-2048G
64-BIT COMMON ALLOCATED = 13292M
```

Where the last row corresponds to the value provided in the SMF71CPA field.

We will enhance EPV to provide information based on SMF71CPA asap.

## Little known SMF parameters



# DDCONS

The DDCONS parameter is included in the SMFPRMxx member of the system parmlib.

It specifies whether duplicate EXCP sections for type 30 SMF records (subtype 2-6) are to be consolidated or not. This is an old parameter created to reduce the number and size of SMF 30 records for address spaces allocating hundreds or even thousands of EXCP sections such as Db2, CICS, IMS, HSM and some long running jobs.

However, the consolidation of these sections may consume a large amount of CPU resources at interval end, step end and especially job end, and take a long time to complete.

The DDCONS default is YES even if all performance analysts, since many years, advice to set DDCONS to NO to save CPU time. In fact, the space saved is negligible compared to the CPU consumed.

The funny thing is that the IBM advice is the same. This is what you can find in the Initialization and Tuning Reference of z/OS 2.5:

“When DDCONS(YES) is specified, SMF merges the EXCP count for these duplicate entries into one entry if the following information is the same:

- DDNAME
- Device class
- Unit type
- Channel address
- Unit address

Long-running jobs might take a long time to end in this case, because of the building of the SMF type 30 records for a long-running job.

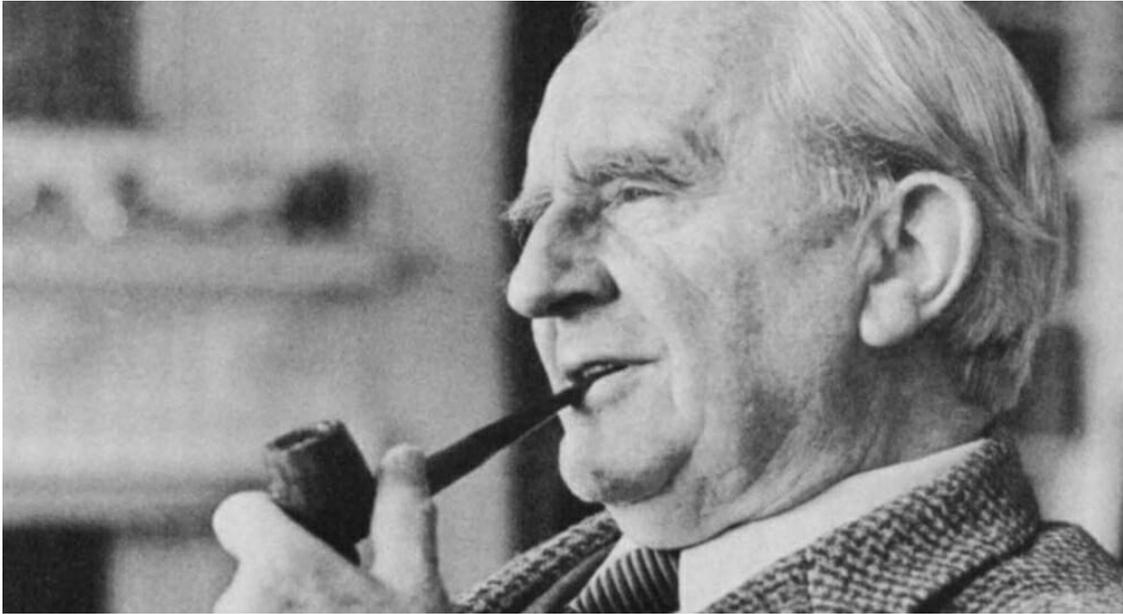
DDCONS(NO) requests that this consolidation function be bypassed, which results in a reduction in the amount of processing required to build the records, and thus a reduction in the amount of time that is required to complete the job.”

Intense consolidation activity could also be reflected in a lower capture ratio.

To understand the impact of DDCONS(YES) you can look at the following fields provided in SMF 30 subtype 4 records (step end):

- SMF30ICU\_STEP\_TERM, CPU (TCB) time used after program end in the previous step;
- SMF30ISB\_STEP\_TERM, CPU (SRB) time used after program end in the previous step.

Setting DDCONS to NO is strongly suggested.



*"You can only come to the morning through the shadows."*

**J. R. R. Tolkien**

---

*Copyright © 2022 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](mailto: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](mailto: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](mailto: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](mailto: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](#)

---

This email was sent to [carlotta.ottaviani@epvtech.com](mailto:carlotta.ottaviani@epvtech.com)  
[why did I get this?](#) [unsubscribe from this list](#) [update subscription preferences](#)  
EPV Technologies · Viale Angelico, 54 · Roma, RM 00195 · Italy

