



epv

IT Cost
Under Control

EPV Technologies

Newsletter

February 2023

THIS MONTH HIGHLIGHTS

- Measuring zFS performance - Part 1
- Important Db2 corrections

Measuring zFS performance - Part 1

The z/OS File System (zFS) is one of the file systems available in z/OS USS (UNIX System Services). However, it's important to note that zFS is the strategic IBM file system and, in general, it provides better performance than the older Hierarchical File System (HFS).

Similarly to HFS, zFS file systems contain files and directories that can be accessed with USS application programming interfaces and can be used for all levels of the USS hierarchy, including the root file system.

zFS usage was initially limited to innovative applications but it is now extended also to the most traditional z/OS workloads such as CICS and IMS.

This is the reason why measuring zFS performance has become more and more important in the last years. Unfortunately, the available documentation about zFS performance analysis and tuning is not very clear and complete yet.

In this paper, after a short introduction to zFS, we will discuss the most relevant parameters which can be set to tune zFS caches.

Then we will look at some powerful zFS commands which can provide a lot of useful performance information.

Finally, we will focus on the available SMF metrics which can be used to measure zFS performance externally, from the z/OS point of view, and internally, from the USS point of view.

If you want to receive the paper you can reply to this e-mail writing "**Measuring zFS performance - Part 1**" in the subject

Important Db2 corrections

PH47163: ENHANCE DB2 REAL STORAGE MANAGEMENT WHEN USING REALSTORAGE_MANAGEMENT=AUTO

There are Db2 users who saw significant increase in Db2 system service address space (SRB) CPU time. This has been observed as one of the effects when processing excessive IARV64 REQUEST(DISCARDDATA) requests.

More details at: [PH47163](#)

PH49442: SLOW DB2 SHUTDOWN DUE TO CHECKPOINT FOR WORKFILE
Db2 subsystem shutdown was delayed by workfiles going through close processing. Close processing for the workfiles was taking a long time because there were many thousands of modified workfile pages that needed to be written out to the workfile datasets during close processing.

More details at: [PH49442](#)



Customer Questions

Hello, I ask you a clarification about the following considerations of my boss on MIPS usage.

“By looking at yesterday EPV pages I see that the total CPU usage of one of our system, in the peak hour, is 14.324 MIPS while the sum of all workloads accounts to 12.824 MIPS.

The difference between these two values should represent the "uncaptured ratio" value.

In this case, it is a bit more than 10% and it looks high to me.

My question is: am I right?

If yes what could be the reasons for that?”

EPV Technical Support answer

Your boss considerations are correct.

A common best practice for production systems is to get more than 90% capture ratio (less than 10% uncaptured). Higher values (up to 95%) are enough common. However, the values also depend on the workload characteristics. Unfortunately, there is not much you can do about that.

Most common factors on which you can act to increase the capture ratio are:

- Paging activity; you can reduce it by giving more memory to systems where paging is significant;
 - I/O activity; you can reduce it by tuning and by exploiting the many available DIM (Data In Memory) possibilities available in z/OS;
 - SRM activity; you can reduce it by checking and eventually increasing the RMPTTOM parameter in OPT.
-



ERV

The ERV (Enqueue Residence Value) parameter is included in the IEAOPTXX member of the system parmlib.

It specifies the number of CPU service units that an address space or enclave is allowed to absorb when it is possibly causing enqueue contention. During this “enqueue residency” time, the address space, including the address space that is associated with an enclave, is not considered for swap-out based on recommendation value analysis. The address space or enclave runs with a high enough priority to guarantee the needed CPU time and hopefully release the enqueue.

The goal is to improve availability and performance of high importance work.

The ERV default is still 500 even if all performance analysts, since many years, advice to set it to 50.000.

IBM agrees with this advice but emphasizes that you should still address the causes of high-importance work delayed by the ENQ contention with lower-importance work.

Our advice is to set ERV to 50.000.

Quotes



"Live life as though nobody is watching, and express yourself as though everyone is listening"

Nelson Mandela

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](#)

This email was sent to 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

