

EPV Technologies

Newsletter

October 2026

THIS MONTH HIGHLIGHTS

- **PHP is not used in EPV products**
- **Saving MSUs with Db2 buffer pools tuning**

PHP is not used in EPV products

We have received inquiries from some clients asking whether the recently discovered vulnerabilities in PHP might affect EPV products.
We can confirm that PHP is not used in any of EPV Technologies' products.

Saving MSUs with Db2 buffer pools tuning

The amount of available mainframe memory continues to increase with every new IBM machine generation as well as the amount of memory supported by the latest z/OS releases. Up to 64 TB memory can be installed in IBM z17 while z/OS 3.2 will support up to 16 TB in a single LPAR.

Large amounts of memory allow the full exploitation of the many available DIM (Data In Memory) techniques.
Maintaining data in memory has two major advantages:

- it improves performance by avoiding the delay caused by I/O operations;
- it eliminates the CPU consumption needed to perform I/O operations.

This is particularly important for workloads that access Db2 buffer pools. Therefore, in this document, we will focus on the potential MSU savings achievable by leveraging the following optimizations offered by current versions of Db2:

- Exploit buffer pool page fixing;
- Exploit buffer pool large frames;
- Reduce buffer pool I/O;
- Use in-memory buffer pools.

If you want to receive this paper, you can reply to this e-mail writing **" Saving MSUs with Db2 buffer pools tuning"** *in the subject.*



Customer question
We were wondering whether you already have any use cases for collecting and reporting on data from IDAA (IBM Db2 Analytics Accelerator for z/OS), either through SMF data or by using the Linux System Activity Report (SAR).
Although we do not yet have an active IDAA environment, we are looking for information and best practices that could help us define our reporting requirements before the solution is implemented and in use.

EPV Technical Support answer
Yes, we have some customers using IDAA.
Of course, we provide IDAA information in EPV for DB2 getting them from standard Db2 SMF records.
If you already use IDAA, we provide reports for each accelerator about:

- CPU UTILIZATION
- DISK UTILIZATION
- REQUESTS
- PERFORMANCE
- THROUGHPUT
- REPLICATION

If you do not use IDAA yet we can provide in EPV for Db2 a rough indication of what applications could save in terms of CPU utilization and elapsed time.
It requires you to activate ACCELERATOR MODELING by setting the ACCELMODEL parm to yes. You also need CACHEDYN=YES to catch dynamic queries; for static queries you need to bind-rebind packages.
In this case in EPV for Db2 WORKLOADS we provide the following reports

- TOTAL ACCEL ELIGIBLE CPU

- SHARING GROUP ACCEL ELIGIBLE CPU
- DB2ID TOP ACCEL ELIGIBLE CPU PLAN
- TOTAL ACCEL ELIGIBLE ELAPSED SHARING GROUP
- ACCEL ELIGIBLE ELAPSED DB2ID
- TOP ACCEL ELIGIBLE ELAPSED PLAN



Advantages and limits of queue statistics

With MQ 9.3 IBM introduced a new class of the statistics trace which provides very relevant queue measurements.

These queue statistics are produced by class 5 of the statistics trace and are written in SMF 115 subtype 216 records.
Class 5 is not active by default so you should activate it by adding it to the appropriate MQ parameters in the SMFSTAT parameter of the CSQ6SYSP macro:

SMFSTAT(1,3,4,**5**)

or by issuing a command:

START TRACE(STAT) CLASS(5)

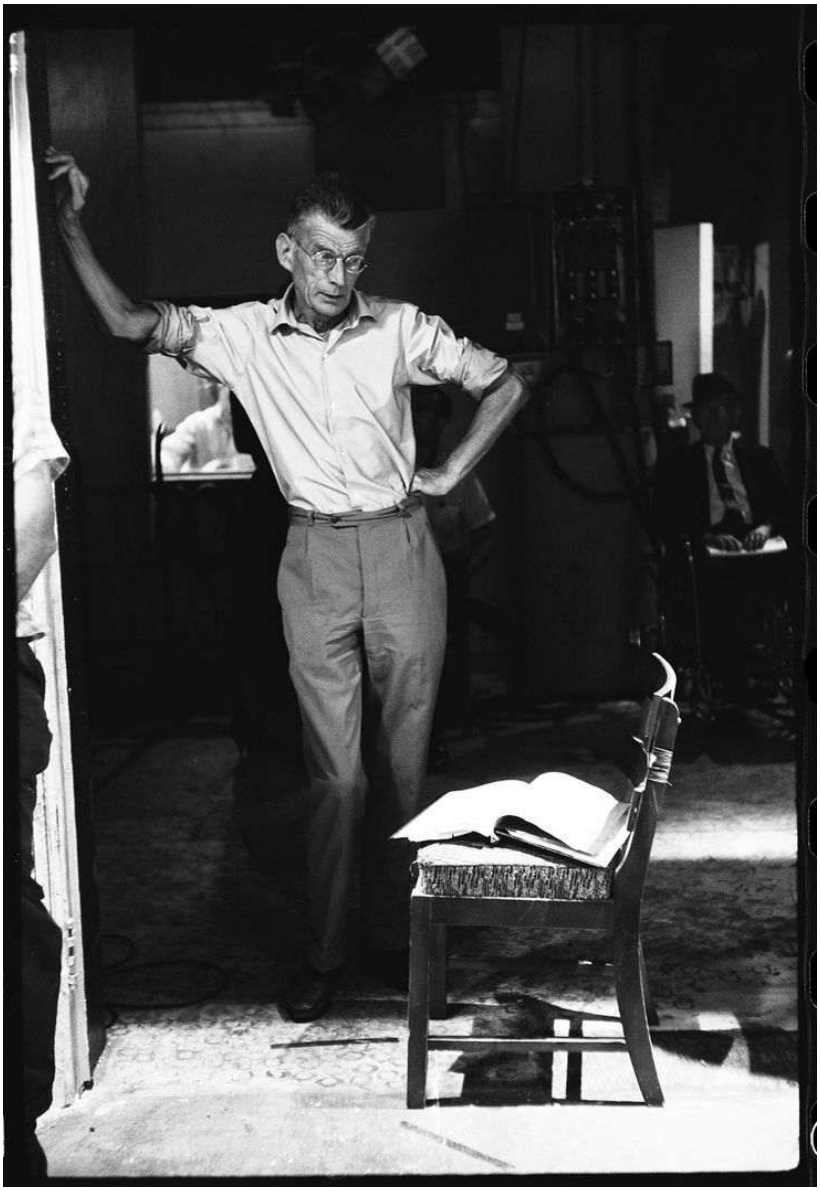
The first option is advisable if you want to collect this information on a regular basis while the second option is normally used for sporadic analysis.
Of course, queues must allow collecting queue statistics with the STATQ(ON) parameter.

These statistics provide a lot of information (more than 80 metrics for each queue), about GETs, PUTs, PUT1s number of messages and size, as well as information about persistent and not persistent messages, latency, queue depth and much more.

The great advantage of queue statistics is that they are collected at the interval set in the STATIME parameter and the number of records depends only on the number of queues. So, the overhead and the number of records produced can be considered as negligible.

The limit of queue statistics compared to queue accounting, which is available from many years, is that there is no information about the workload accessing the queues.
If you need this information, the only solution is the activation of the class 3 of the accounting trace which produces SMF 116 subtype 2 and 3 records. In this case due to the introduction of information about threads accessing each queue, the overhead and the number of records produced can be significant.

Quotes



"Ever tried. Ever failed.
No matter.Try Again.
Fail again. Fail better"
Samuel Beckett

We care about your Privacy. EPV Technologies is GDPR-compliant.

You may have heard about the new General Data Protection Regulation ("GDPR"), that comes into effect May 25, 2018.
It was introduced to unify all EU countries to a unique data regulation, ensuring that all data protection laws are applied identically within the EU.
It also protects EU citizens from organisations using their data irresponsibly and puts them in charge of "what", "where" and "how" information is shared.

To see our Privacy Policy click here
EPV Technologies Privacy Policy

Your continued subscription is considered acceptance
of the Terms and Conditions placed on the following link:
EPV Technologies Terms and Conditions

Copyright © 2026 EPV Technologies, All rights reserved.
You have the right to remove yourself from the newsletter subscription list at any time. If at any time you wish to unsubscribe, there is a link at the bottom of this email, or any subsequent newsletter you receive. You can also unsubscribe by simply sending a mail to epv.info@epvtech.com with the subject "REMOVE FROM TECHNICAL NEWSLETTER ".

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


INTUIT
mailchimp