



**epv**

IT Cost  
Under Control

# EPV Technologies

## Newsletter

May 2021

### THIS MONTH HIGHLIGHTS

- Easy and Cheap Compression with EPVzip V15
- EPV Performance University 2021 - Second Edition
- IBM Unveils World's First 2 Nanometer Chip Technology

### Easy and Cheap Compression with EPVzip V15

The policy of EPV products has always been to provide all the needed information for z/OS performance analysis and capacity planning, by running on an external platform such as Windows, Linux or Unix.

The reason for this choice is our continuous effort to reduce customer hardware and software costs.

The main source of information for all the capacity management activities are:

- SMF records, which are in VBS format,
- IMS log records and DCOLLECT records, which are in VB format.

When transferring this data to an external platform there are two important issues to consider:

- the amount of data to be transferred is normally very large,
- you need to preserve data integrity, control blocks included, in order to be able to correctly parse them.

Both issues can be addressed with compression. This is the reason why, some years ago, we developed the EPVzip tool which, running on z/OS, can compress VBS and VB data before transferring them to the processing platform.

Of course, we wanted to limit as much as possible the impact on z/OS costs, so we developed it in Java to allow it running on zIIP instead of CPU.

The result is very good: at least 90% of the EPVzip consumptions are on zIIP and the compression ratio is normally around 10.

Even if SMF data are already compressed, as it could be the case for the SMF records produced by CICS and Db2, EPVzip provides a meaningful reduction of the data size.

It is important to note that EPVzip uses the zlib compression library, so, when running on z15 machines, it automatically exploits the Integrated Accelerator for zEDC.

In this paper we will show what you need to do to use EPVzip.

We will also discuss EPVunzip: an additional tool which allows users to decompress, on z/OS, any VBS and VB file, previously compressed using EPVzip.

If you want to receive the paper you can reply to this e-mail writing **"Easy and Cheap Compression with EPVzip V15"** in the subject

---

**EPV Performance University 2021**  
**Second Edition**

---

If you have lost the Winter edition this is an opportunity you don't have to miss.

The EPV Performance University will be virtual.

It will be free of charge and reserved to EPV customers, partners and invited guests.

It will be held on two weeks: June 8-10, 2021 and June 15-17 2021.

Mark these days in your agenda.

The EPV Performance University will be structured in three training courses on the following topics:

- z/OS performance analysis, (3 days)
- WLM update (2 days)
- Db2 performance analysis (1 day)

These training courses have been designed to provide participants a deeper knowledge about:

- the most relevant performance metrics to use to analyse z/OS performance;
- the most advanced and recent WLM functions;
- the most relevant performance metrics to use to analyse Db2 performance.

Here is the agenda

## Week 1, 8-10 June 2021

### z/OS Performance Analysis - Day 1

Date	Start	End	Description
08/06/2021	09:00	10:00	z/OS data collection - Part 1
08/06/2021	10:15	11:15	z/OS data collection - Part 2
08/06/2021	11:30	12:30	CPU
08/06/2021	14:30	15:30	WLC and TFP
08/06/2021	15:45	16:45	zIIP

### z/OS Performance Analysis - Day 2

Date	Start	End	Description
09/06/2021	09:00	10:00	SMF 113 counters
09/06/2021	10:15	11:15	Memory
09/06/2021	11:30	12:30	Coupling Facility
09/06/2021	14:30	15:30	Disk I/O
09/06/2021	15:45	16:45	Crypto

### z/OS Performance Analysis - Day 3

Date	Start	End	Description
10/06/2021	09:00	10:00	Address spaces
10/06/2021	10:15	11:15	Service and Report classes
10/06/2021	11:30	12:30	CICS and IMS throughput
10/06/2021	14:30	15:30	Other workloads throughput
10/06/2021	15:45	16:45	Daily Trends

## Week 2, 15-17 June 2021

### WLM Update - Day 1

Date	Start	End	Description
15/06/2021	09:00	10:00	WLM basics - Part1
15/06/2021	10:15	11:15	WLM basics - Part2
15/06/2021	11:30	12:15	WLM basics - Part3
15/06/2021	14:30	15:15	SYSTEM, SYSSTC and protection
15/06/2021	15:30	16:15	CICS and IMS transaction goals

### WLM Update - Day 2

Date	Start	End	Description
16/06/2021	09:00	10:00	Capping and soft capping
16/06/2021	10:15	11:15	Resource Groups
16/06/2021	11:30	12:15	IRD and HiperDispatch
16/06/2021	14:30	15:15	zIIP and SMT
16/06/2021	15:30	16:15	WLM and software pricing

### Db2 Performance Analysis

Date	Start	End	Description
17/06/2021	09:00	10:00	Data collection
17/06/2021	10:15	11:15	System AS
17/06/2021	11:30	12:15	EDM pools and threads
17/06/2021	14:00	14:45	Buffer Pools and Group Buffer Pools
17/06/2021	15:00	15:45	Locks and latches
17/06/2021	16:00	17:00	Workload analysis

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.

The courses language will be English

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

---

## **IBM Unveils World's First 2 Nanometer Chip Technology**

---

ALBANY, N.Y., May 6, 2021 /PRNewswire/ -- IBM (NYSE: IBM) today unveiled a breakthrough in semiconductor design and process with the development of the world's first chip announced with 2 nanometer (nm) nanosheet technology. Semiconductors play critical roles in everything from computing, to appliances, to communication devices, transportation systems, and critical infrastructure...

The potential benefits of these advanced 2 nm chips could include:

- Quadrupling cell phone battery life, only requiring users to charge their devices every four days.
- Slashing the carbon footprint of data centers, which account for one percent of global energy use. Changing all of their servers to 2 nm-based processors could potentially reduce that number significantly.
- Drastically speeding up a laptop's functions, ranging from quicker processing in applications, to assisting in language translation more easily, to faster internet access.
- Contributing to faster object detection and reaction time in autonomous vehicles like self-driving cars.

More information at:

[IBM Unveils World's First 2 Nanometer Chip Technology](#)

---



## Customer Questions

We had problems with ECSA utilization. It got saturated and caused a crash of our development system. We need to enlarge it by reducing the extended private area accordingly, but we want to avoid hurting some of our address spaces.

Is there a way to know how much extended private storage is used by each address space?


### ***EPV Technical Support answer***

Extended private storage (above the 16MB line) is composed by the following areas:

- ELSQA, Extended Local System Queue Area, it contains tables and control blocks queues associated with the address space;
- ESWA, Extended Scheduler Work Area; it contains control blocks that exist from task initiation to task termination; it includes control blocks and tables created during job control language (JCL) interpretation;
- USER, this area is used for running user program applications and storing user program data.

Their usage is measured in the following fields of SMF 30 records:

- SMF30EAR, maximum virtual storage in bytes allocated from the ELSQA and the ESWA subpools;
- SMF30EUR, maximum virtual storage in bytes allocated from the user subpools.



# Major suspension reasons for Db2 applications

## Log Write I/O suspensions

The time needed for synchronous log write I/Os, when they do not occur under an Execution Unit (EU) Switch, is accounted as Log Write I/O suspension time. This time directly impact the wait and elapsed times of Db2 applications.

When logging to disk occurs under an EU Switch (e.g. phase 2 commit processing for updates, inserts, and deletes), that time is accounted only in a different suspension reason to avoid double accounting.

In general, Log Write I/O suspensions occur:

- at commit time: 2 synchronous log writes for two-phase and 1 synchronous log write for one-phase commit;
- at a synchronous write to disk or GBP; it is driven by the Log Write Ahead (LWA) protocol which dictates to write to the log first;
- before starting any undo processing (e.g. rollbacks);
- when running out of log buffers; it triggers writing to the active log;
- when the immediate write threshold (97,5% of unavailable pages) is reached for a Buffer Pool.

When in data sharing, they also occur because of:

- Page p-lock negotiation; when a changed page has to be written to a GBP as part of the negotiation; LWA protocol;
- Index page split of a GBP dependent index;
- Index page deletion of a GBP dependent index;
- Setting the IMMEDIATE bind option or the IMMEDIATE subsystem parameter to YES; LWA protocol;
- Using GBPCACHE(NO); changed pages written to disk at commit; LWA protocol;
- An overflow record is added for a GBP dependent object;
- Space map page update for a GBP dependent object;
- IFI READS IFCID 306 request from another member in the data sharing group.

At plan level (IFCID 003), the following metrics are provided in the following fields of SMF 101 records:


- QWACAWLG, accumulated wait time for log write I/O;
- QWACARLG, number of log I/O suspensions.

If the average I/O time is too long, disk I/O performance should be investigated.

A best practice can be less than 2 milliseconds, but it depends on many factors such as:

- used storage processors,
- design of the I/O infrastructure,
- distance between systems and disks,
- synchronous remote copy,
- overlapping of log write I/O to the primary and secondary disk subsystem in a synchronous peer-to-peer remote copy environment (zHyperWrite technology),
- exploitation of synchronous I/Os (zHyperLink technology).

## Quotes



*"Just because something doesn't do what you planned it to do doesn't mean it's useless.*

**Thomas Edison**



---

*Copyright © 2021 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



