EPV Technologies

Newsletter

February 2021

- Measuring the Integrated Accelerator for zEDC
- EPV Performance University 2021 – Second week is about to start
- Withdrawal of Country Multiplex Pricing offering

Measuring the Integrated Accelerator for zEDC

In the last years we have seen an explosion of data, driven by new technologies and applications. This is true for every environment, also for z/OS systems.

Even if storage subsystems have become more and more powerful, they can only solve part of the problem. All this data has to be managed by system components and applications with a consequent increase of elapsed time and CPU consumptions.

A more complete solution is compression. Compression may greatly reduce data size with great benefits in terms of storage usage and performance.
Even if software compression has been available in various forms for many years, it has only partially been exploited by customers because of the CPU cycles needed to execute compression/decompression instructions.

To reduce the amount of CPU needed, with the zEC12 machine IBM introduced zEnterprise Data Compression (zEDC) a hardware/software solution based on a specialized adapter for data compression, known as zEDC Express.

With the z15 machine, IBM eliminated the need of the zEDC Express cards by introducing the Integrated Accelerator for zEDC (accelerator in the following). This is a much more powerful solution which could also be the key to allowing a real exploitation of pervasive encryption. In fact, another benefit of compression is that it can greatly reduce the amount of data to encrypt.

In this paper, after a short introduction to the accelerator, we will discuss the characteristics of the synchronous and asynchronous compression requests both in terms of exploiters and software licenses.

Finally, we will also discuss which SMF metrics are available to measure these two types of compression requests showing some real-life examples.

If you want to receive the paper you can reply to this e-mail writing **"Measuring the Integrated Accelerator for zEDC"** in the subject.

EPV Performance University 2021
Second week is about to start

The first week of the virtual edition of the EPV Performance University has just ended.
Both the WLM Update and Db2 Performance Analysis training courses have been very successful with an average of 65 participants from 30 different companies every day.

We take this opportunity to thank all the participants.

The second week is about to start.

You still have time to subscribe to the z/OS Performance Analysis training course that will be held on 2-4 March.
Participation is free of charge for EPV customers, partners and invited guests.

The courses language will be English

All details, final agenda and subscription forms are available at:
www.epvtech.com

Withdrawal of Country Multiplex Pricing offering

Effective January 1, 2021, IBM will withdraw the Country Multiplex Pricing (CMP) offering announced in Software Announcement 215-230, dated July 28, 2015.

CMP is a sub-capacity offering that allows clients to use their IBM Z® server capacity within a given country without the constraints of sysplex aggregation rules and many of the limitations of previous reporting methodologies.
This offering is being replaced by the IBM Z strategic pricing option, Tailored Fit Pricing for IBM Z software, announced in Software Announcement 219-014, dated May 14, 2019.

Tailored Fit Pricing for IBM Z is a transformational pricing option for IBM Z software. It offers simple, transparent, and predictable pricing for IBM Z software running on the IBM z/OS platform within a given country.

This announcement does not affect existing CMP clients with standard, non-expiring contract terms. These clients will continue to operate under CMP terms and conditions, including those CMP clients who have transitioned to a Tailored Fit Pricing for IBM Z Enterprise Solution…

Full announcement letter at:
[Withdrawal of Country Multiplex Pricing offering - Replacement available](#)

Customer Questions

How can I calculate the amount of MSUs used by a system?
And how the MSUs used by a single address space?

EPV Technical Support answer
The algorithm is always the same:
a)   you need to divide the amount of CPU seconds used by the seconds in 1 hour (3600); this will provide you the number of CPUs used
b)   you need to divide the machine capacity in MSUs, provided by the SMF70MCR field in the SMF 70 records, by the number of CPUs in the machine; this will provide you the MSU capacity of a single CPU
c)   Finally, you need to multiply the results of a) and b)

MSUs used = (CPU used / 3600) * (SMF70MCR / #CPUs in CEC)

If you apply this formula to address spaces, service classes or report classes you will underestimate the MSUs used. To have a better estimate you need to divide the MSUs used by the System Capture Ratio.

Similar considerations (a bit more complex) apply to the MSUs used by CICS and IMS transactions.
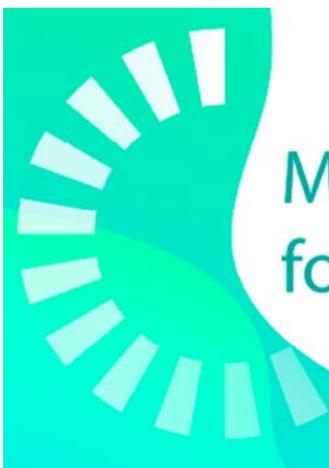
Major suspension reasons for Db2 applications

Lock/Latch suspensions (part 2)

Lock/latch suspensions include both IRLM local lock and IRLM latch suspensions. The suspension time is provided in the QWACAWTL SMF metric.

To have an indication on the amount of each suspension type you can look at the QTXASLOC (number of IRLM local lock suspensions) and QTXASLAT (number of IRLM latch suspensions) SMF metrics.
Unfortunately, they do not tell you anything about the suspension time for each type. Generally speaking, IRLM latch suspensions should be shorter than lock suspensions.

IRLM local lock suspensions

IRLM manage lock requests in chains and serialize the access to these chains by using its specific latch type, which is called IRLM latch
IRLM latch suspensions happen when there is contention on the same latch chain.
A best practice is that IRLM latch contention should not be higher than 5% (IRLM latch suspensions should represent no more than 5% of all the IRLM lock requests).

To calculate IRLM latch contention at the subsystem level you can use the metrics available in SMF 100 records and the following formula:

QTXASLAT / (QTXALOCK + QTXAUNLK + QTXAQRY + QTXACHG)

where:
* QTXASLAT, IRLM latch suspensions
* QTXALOCK, lock requests
* QTXAUNLK, unlock requests
* QTXAQRY, lock query requests
* QTXACHG, lock change requests

By using the same formula, with the metrics provided in SMF 101 records, IRLM latch contention can be calculated also at the application level.

Most common reasons of high IRLM latch contention and suggested actions are:
* IRLM trace on; turn it off;
* IRLM dispatching priority too low; IRLM has to run in SYSSTC;
* Frequent IRLM commands (e.g. DISPLAY DATABASE LOCKS); check automation and monitoring tools;
* Deadlock detection cycle low and high locking rate; check if you can increase the detection cycle or improve lock avoidance;
* Many RELEASE(DEALLOCATE) packages are used; optimizations have been

introduced in recent Db2 versions to reduce the impact on IRLM latch suspensions; if you are at Db2 10 you could move to Db2 11 or 12.

Quotes

The cleverest of all, in my opinion, is the man who calls himself a fool at least once a month.

Fyodor Dostoevsky

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