



epv

IT Cost
Under Control

EPV Technologies

Newsletter

September 2026

THIS MONTH HIGHLIGHTS

- **Saving the MSUs used by zIIP-eligible workloads**
- **EPV User Group 2026 final agenda and abstracts**

Saving the MSUs used by zIIP-eligible workloads

In 2006 IBM released the zIIP (z Integrated Information Processor), a specialty processor initially dedicated to run enclave SRB work.

In the years since, many other workloads have been made completely or partially zIIP-eligible; the most relevant are:

- any Java-based application;
- DDF requests from external network clients via TCP/IP querying the local Db2 on z/OS;
- most of DB2 DBM1 activities;
- Db2 MSTR log activities;
- Db2 parallel child tasks generated by complex SQL queries (e.g., Data Warehousing or Business Intelligence);
- Db2 utilities and statistics gathering;
- Python-based applications;
- /OS XML System Services for parsing large XML files;
- Docker containers and Linux workloads running in z/OS Container Extensions (zCX).

However, eligibility is a necessary but not sufficient condition. A zIIP engine must be available with an acceptable delay, according to the “zIIP needs help” mechanism when the workload request arrives.

Fully leveraging zIIP technology is a key element of the strategy for any customer aiming to reduce hardware costs and, above all, software costs especially when the TFP Software Consumption Solution pricing option has been adopted and every MSU counts.

The goal must be to minimize the portion of zIIP-eligible workload that spills over to standard CPUs because a zIIP engine was not available at the time.

In this document, following a brief description of the "zIIP needs help" mechanism, we will discuss zIIP capacity planning, the Simultaneous Multi-Threading (SMT) technology, and the specific parameters you can configure to achieve this goal.

*If you want to receive this paper, you can reply to this e-mail writing **"Saving the MSUs used by zIIP-eligible workloads"** in the subject.*

EPV User Group 2026 final agenda and abstracts

The 24th EPV User Group will once again take place virtually. The event will be held on October 27 and repeated on October 29.

This year, too, we will kick off the day with two training sessions dedicated to new Performance Analysts.

The EPV User Group is a "not to miss" event for all Performance Analysts; it will give you the opportunity to share ideas with qualified experts and to listen to some of the EPV customers experiences. The most interesting features provided by the latest versions of all EPV products will also be presented.

Agenda			
27 and 29 October 2026			
10:00 – 10:10	Welcome and introduction	Danilo Gipponi	EPV Technologies
10:10 – 10:30	PRSM overhead	Fabio Massimo Ottaviani	EPV Technologies
10:30 – 10:50	Capture ratio	Danilo Gipponi	EPV Technologies
10:50 – 11:10	EPV zParser and EPV Real Time V18	Matteo Bottazzi	EPV Technologies
11:10 – 11:30	coffee break		
11:30 – 11:50	EPV zParser: Parsing SMF data into Elasticsearch	Imri Keidar	FIBI
11:50 – 12:10	EPV for CICS V18 preview	Luca Battaglio	EPV Technologies
12:10 – 12:30	Interactive analysis of individual CICS transactions and Db2 Packages with MyEPV	Stefano Rotunno	EPV Technologies
12:30 - 14:00	Lunch		
14:00 – 14:20	EPV for MQ V18 preview	Luca Trazzi, Fabio Massimo Ottaviani	EPV Technologies
14:20 – 14:40	A reporting system methodology to keep your CPU consumptions under control	Maurizio Ciucci	BCC Sistemi Informatici
14:40 – 15:00	EPV for z/OS V18	Massimo Orlando	EPV Technologies
15:00 – 15:20	coffee break		
15:20 – 15:30	zIIP on CP (IIPCP), possible causes and effects	Hans Gerd Schneider	BarmeniaGothaer
15:30 – 16:00	IBM z17 performance results	Fabio Massimo Ottaviani	EPV Technologies

	New Performance Analysts Education
	User presentation
	Product presentation
	Technical update

Abstracts

PRSM overhead (AE)

In this tutorial, we will examine the metrics and reports that can be used to monitor PR/SM overhead. We will also suggest reference thresholds and discuss possible causes of excessive overhead. Finally, we will explain the Hiper Dispatch mechanism, the most powerful tool available for reducing PR/SM overhead.

Capture ratio (AE)

The "capture ratio" is an inverse measure of system overhead. It represents the proportion of total CPU or zIIP time that z/OS is able to attribute to specific address spaces or enclaves. In this tutorial, we will outline widely recognized best practices, as well as the metrics and reports used to monitor the capture ratio. We will also examine the possible causes of a low capture ratio and the actions that can be taken to improve it.

EPV zParser and EPV Real Time V18 (PP)

EPV zParser and EPV Real Time share some of their code.

EPV zParser is designed to read files (including SMF and other record types) transferred from z/OS to an external platform, parse all the records, and store them in a database; it forms the foundation of all major EPV products (EPV for z/OS, EPV for Db2, etc.).

EPV Real Time is designed to read SMF records, collected in an in-memory resource, from an external platform via TCP/IP and store them in a database; it forms the foundation of EPV Real Time CICS and EPV Real Time STC.

In this presentation, we will outline the most significant improvements introduced in this new version of both products.

EPV zParser: Parsing SMF data into Elasticsearch (UE)

This presentation covers a project analyzing SMF data from CDC/IIDR (IBM InfoSphere Data Replication) to measure CPU consumption per Subscription and schema.

Data flows from the Mainframe (SMF) through an EPV server on Windows NT, then to a Linux server that streams the data into Elasticsearch (including Python code using JSON and requests).

We will show the results of analyzing CPU consumption trends over time — by Subscription and by schema.

EPV for CICS V18 preview (PP)

In this preview, we will examine the most significant views and features being introduced in version 18. These include a brand new TCP view and additional views related to transaction abends and RNI activity. New critical events will be added, and, even more importantly, daily, weekly, and monthly trends for these events will be introduced.

Interactive analysis of individual CICS transactions and Db2 Packages with MyEPV (UE)

Locating performance problems for individual CICS transactions is a very complicated task in big environments executing millions of transactions a day and connecting to various subsystems such as Db2, MQ, WebSphere and others. In this presentation we will show you how we succeeded to solve this problem using the zParser and the MyEPV Dynamic reporting tool.

EPV for MQ V18 preview (PP)

In this preview, we will present some of the enhancements being introduced in EPV for MQ V18. These include new views, based on queue statistics, to monitor average message residence time, average and total message size, as well as failures and rollbacks for MQGET and MQPUT operations. We will also add an efficiency index for MQGET operations and an estimate of polling time.

A reporting system methodology to keep your CPU consumptions under control (UE)

Many times, we are surprised by significant CPU increases that lead to higher costs or performance issues. Given the volume of z/OS SMF data and the variability of workload CPU consumption, it is difficult to detect such increases in a timely manner before they escalate into a problem. In this presentation, we will demonstrate a reporting system that enables easy detection of anomalous behavior and identification of the root cause using various techniques that will be described.

EPV for z/OS V18 (PP)

The new version of EPV for z/OS, along with EPV zParser, is always the first, in the EPV product suite, to be developed. Consequently, the presentation will highlight the most significant improvements in usability and manageability introduced in all EPV V18 products. These include support for z17, a redesigned exception management system, EPV Report Assistant, and EPV Knowledge Center. Key new views regarding DPU usage, address space memory occupancy, and statistical monitoring of batch jobs will also be presented.

zIIP on CP (IIPCP), possible causes and effects (UE)

This presentation focuses on the general approach adopted by the supervisor/dispatcher, based on parameter settings and methods for assigning zIIP-eligible workloads to GCPs.

The presentation also examines some surprising effects and the associated analysis.

IBM z17 Performance Results (TU)

We regularly receive requests from clients to evaluate the performance of a new machine; the question is always the same: "Is my new machine delivering the expected performance?". In this presentation, we propose a methodology to address this question by defining and verifying expectations and selecting the most significant data and metrics to compare the performance of the old and new machines. This methodology has been applied to a real-world case study that serves as the common thread throughout the document.

[Subscription form](#)

The EPV User Group is a "not to miss" event for all Performance Analysts; it will give you the opportunity to share ideas with qualified experts and to listen to some of the EPV customers experiences. The most interesting features provided by the latest versions of all EPV products will also be presented.

The EPV User Group is free of charge and reserved to EPV customers. If you are not a customer yet but you are interested in participating, please answer to this e-mail asking for an invitation.

Mark these dates in your agenda to avoid missing this event.

You can subscribe at: www.epvtech.com



Customer question

We wanted to calculate the number of requests for each z/OS Connect API. Should we count the records for a specific time slot? Or is there a field we can derive it from?

EPV Technical Support answer

You need to decide which timestamp to use to set the date and time for aggregation (e.g., date, hour, minute) and then count the number of records (requests) for each API.

If we consider the case where z/OS exposes the APIs (subtype 1, version 2), we have this time measurements:

- SMF123S1_TIME_ZC_ENTRY, "Time the request was received by the z/OS Connect EE server."
- SMF123S1_TIME_ZC_EXIT, "Time the response was ready to be sent to the HTTP client."
- SMF123S1_TIME_SOR_SENT, "Time the request was sent to the system of record."
- SMF123S1_TIME_SOR_RECV, "Time the response was received from the system of record."

We suggest using SMF123S1_TIME_ZC_EXIT.



Channel initiator CPU usage

The MQ Channel Initiator (CHIN) address space may use a significant amount of CPU at many sites.

While SMF 30 and 72 allows you just to measure it as a black box, by activating class 4 of the statistics trace you can get interesting details on CHIN internal processing.

These statistics are collected in subtype 231 of the SMF 115 at the interval set in the STATIME parameter; the overhead and the number of records produced can be considered as negligible.


With this information you can decompose the CHIN CPU usage in the following components

- a. dispatcher tasks; they manage the communication between channels and adapters; by default, 5 tasks are activated but customers can change their number by customizing the CHIDISPS parameter;
- b. adapters tasks; they manage the communication between dispatchers and queue manager; by default, 8 tasks are activated but customers can change their number by customizing the CHIADAPS parameter;
- c. SSL tasks; they manage the encryption/decryption process to allow the queue manager to communicate securely with other queue manager or clients; SSL tasks are optional; The number of tasks to activate must be set in the SSLTASKS parameter;
- d. DNS task; there is only one DNS task; it is used when a channel starts, to convert an output connection name to an IP address, and from an input IP address to a connection name.

In the record section of each component the QCTCPTM provides the total CPU time used.

Additional metrics about throughput and elapsed time are also provided.

The following figure shows an example of a graph that allows you to examine the CHIN CPU usage within the address space.



You can see that dispatcher tasks use most of the CPU, followed by adapters. SSL tasks CPU usage is also significant while DNS tasks usage is negligible.

You can also note that SMF 115-231 accounting is good. More than 90% of the CPU used by the address space is attributed to CHIN components.

Quotes



“Learning never exhausts the mind.”
Leonardo Da Vinci

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
Via Lugi Mancinelli, 106
Roma, RM 00199
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](#)

