

EPV User Group 2026

Performance and Capacity Planning
for Mainframe Environments

27 and 29 October 2026
Virtual Conference

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	to be confirmed	to be confirmed
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



epv

IT Cost
Under Control

EPV customers can register
filling up the form at
www.epvtech.com

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 increasements 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.