



epv

IT Cost
Under Control

EPV Technologies

Newsletter

September 2022

THIS MONTH HIGHLIGHTS

- Monitoring connections security with zERT - Part 1
- EPV User Group 2022 – Final Agenda with Abstracts
- IBM Redbooks - Optimized Inferencing and Integration with AI on IBM Z

Monitoring connections security with zERT Part 1

Security is one of the most important strengths of the z/OS system. This is the reason why IBM continues to enhance data security with initiatives such as pervasive encryption.

Even if not all companies require pervasive encryption, most companies need to encrypt network traffic between the mainframe and the outside world.

With z/OS 2.3 IBM introduced the z/OS Encryption Readiness Technology (zERT). It is a new capability of the Communications Server which allows the

collection of information from the TCP/IP stack about cryptographic security attributes of IPv4 and IPv6 application traffic, protected using the TLS, SSL, SSH and IPsec cryptographic network security protocols. Information about Enterprise Extender connections are also collected.

This information can be written to new subtypes of SMF 119 records and analyzed to improve connections security.

In this paper we will explain how to customize TCP/IP to collect zERT information and which are the most relevant measurements collected. We will also show some examples of reports which can be used to analyze and improve the security of your connections.

If you want to receive the paper you can reply to this e-mail writing **"Monitoring connections security with zERT - Part 1"** in the subject

EPV User Group 2022 Final Agenda with Abstracts

EPV User Group 2022 Virtual Edition

24-27 October 2022



Performance and Capacity Planning
for mainframe environments

The XX EPV User Group will be a "virtual" user group. To allow for the widest possible participation, all sessions will be repeated twice and spread across four days from 24th to 27th October.

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.

Session A						
A1	24/10/2022	Monday	09:30 – 09:45	Introduction - First day	Danilo Gipponi	EPV Technologies
	24/10/2022	Monday	09:45 – 10:15	Latest enhancements in EPV for zOS V16	Massimo Orlando	EPV Technologies
	24/10/2022	Monday	10:30 – 11:00	Are you ready for TFP?	Mark Cohen Austrowiek	EPV Technologies
	24/10/2022	Monday	11:15 – 11:45	From z14 to z16: a user experience	Danilo Gipponi	EPV Technologies
	24/10/2022	Monday	14:00 – 14:15	10 management questions you can answer in seconds with EPV Graph for z/OS	Fabio Massimo Ottaviani	EPV Technologies
	24/10/2022	Monday	14:30 – 15:00	Using EPV Real Time to solve user-oriented monitoring needs	Roberto Gioi	MPS
	24/10/2022	Monday	15:15 – 15:45	z16 Capacity Planning	Danilo Gipponi	EPV Technologies
A2	25/10/2022	Tuesday	09:30 – 09:45	Introduction - Second day	Danilo Gipponi	EPV Technologies
	25/10/2022	Tuesday	09:45 – 10:15	Exploiting zERT SMF Data with EPV zParser	Michele Giannuli, Stefano Sidoli	CSE
	25/10/2022	Tuesday	10:30 – 11:00	Ad-hoc analysis with MyEPV V16	Mark Cohen Austrowiek	EPV Technologies
	25/10/2022	Tuesday	11:15 – 11:45	Work Less and Gain More 2.0	Yosh Kasmirsky	Discount Bank
	25/10/2022	Tuesday	14:00 – 14:15	Feeding a data lake with EPV zParser	Matteo Bottazzi	EPV Technologies
	25/10/2022	Tuesday	14:30 – 15:00	What's new in EPV for Db2 V16	Dana Cohen Austrowiek	EPV Technologies
	25/10/2022	Tuesday	15:15 – 15:45	New measurements in Db2 V13	Fabio Massimo Ottaviani	EPV Technologies
Session B						
B1	26/10/2022	Wednesday	09:30 – 09:45	Introduction - First day	Danilo Gipponi	EPV Technologies
	26/10/2022	Wednesday	09:45 – 10:15	Latest enhancements in EPV for zOS V16	Massimo Orlando	EPV Technologies
	26/10/2022	Wednesday	10:30 – 11:00	Are you ready for TFP?	Mark Cohen Austrowiek	EPV Technologies
	26/10/2022	Wednesday	11:15 – 11:45	From z14 to z16: a user experience	Danilo Gipponi	EPV Technologies
	26/10/2022	Wednesday	14:00 – 14:15	10 management questions you can answer in seconds with EPV Graph for z/OS	Fabio Massimo Ottaviani	EPV Technologies
	26/10/2022	Wednesday	14:30 – 15:00	Using EPV Real Time to solve user-oriented monitoring needs	Roberto Gioi	MPS
	26/10/2022	Wednesday	15:15 – 15:45	z16 Capacity Planning	Danilo Gipponi	EPV Technologies
B2	27/10/2022	Thursday	09:30 – 09:45	Introduction - Second day	Danilo Gipponi	EPV Technologies
	27/10/2022	Thursday	09:45 – 10:15	Exploiting zERT SMF Data with EPV zParser	Michele Giannuli, Stefano Sidoli	CSE
	27/10/2022	Thursday	10:30 – 11:00	Ad-hoc analysis with MyEPV V16	Mark Cohen Austrowiek	EPV Technologies
	27/10/2022	Thursday	11:15 – 11:45	Work Less and Gain More 2.0	Yosh Kasmirsky	Discount Bank
	27/10/2022	Thursday	14:00 – 14:15	Feeding a data lake with EPV zParser	Matteo Bottazzi	EPV Technologies
	27/10/2022	Thursday	14:30 – 15:00	What's new in EPV for Db2 V16	Dana Cohen Austrowiek	EPV Technologies
	27/10/2022	Thursday	15:15 – 15:45	New measurements in Db2 V13	Fabio Massimo Ottaviani	EPV Technologies

Abstracts

Latest enhancements in EPV for z/OS V16 (PP)

This presentation will describe the latest enhancements introduced in EPV for z/OS V16.

Part of them relate to IBM z16 support while some others come from our customers experience or have been suggested by them.

Are you ready for TFP? (PP)

Being ready for TFP means you need to be able to use all the IBM MSU's you have committed to in your IBM TFP contract and manage your service levels efficiently.

In this presentation we will show you how EPV makes all these tasks reliable and easy for you.

From z14 to z16: a user experience (UE)

When moving to a new machine you normally have great expectations. Especially if you are doing a two-generation jump. But what you need to measure to understand if your expectations are met? We will discuss that in this presentation.

10 management questions you can answer in seconds with EPV Graph for z/OS (PP)

There are questions most technical managers ask periodically. They normally

want a quick answer and they like graphs.

In this presentation we will show how to find the right answers to the most frequent management questions with EPV Graph for z/OS.

Using EPV Real Time to solve user-oriented monitoring needs (UE)

Every installation has its own unique monitoring issues so flexibility is mandatory for tech software.

In this session it is shown how EPV RT capabilities can provide strong support to fine-grained special monitoring topics as well as answers to particular top management requests.

z16 Capacity Planning (MT)

In this presentation, we'll have a look at the most important capacity characteristics of the IBM z16.

Then, we'll compare z15 and z16 processor cache architecture. Finally, we'll analyse in more detail the new z16 Measurement Facility basic and extended counters provided in SMF 113, using them to calculate the most important indexes to use in performance analysis and capacity planning.

Exploiting zERT SMF Data with EPV zParser (UE)

zERT detail and summary SMF records make available a lot of useful information about cryptographic protection of TCP connections and Enterprise Extender (EE) connections. In this session we show how we have extracted vital information related to certificates, used in our data center, using EPV zParser. We also provide other examples of possible use of these records.

Ad-hoc analysis with MyEPV V16 (PP)

The new MyEPV GUI v16 version provides many new functionalities that will make your ad-hoc performance analysis and the process for scheduling reports much easier.

In this presentation we will show you all these new enhancements.

Work Less and Gain More 2.0 (UE)

In this presentation we will present a new capability of MyEPV/Quick View, enabling a fully automated comparison of resource consumption between similar days. For example, if the compared day is attributed as a "Red day" and we want to compare it against days in the previous quarter, then the comparison will only be made against those days that are attributed as "Red days" and not against other days.

Feeding a data lake with EPV zParser (MT)

An important feature of EPV zParser is the possibility to integrate with the most popular data lakes. In this presentation we will show some of the possible methods you can use to feed Hadoop, Splunk and Elasticsearch starting from EPV zParser data.

What's new in EPV for Db2 V16 (PP)

The new EPV for Db2 version is in the final development stages. In this

presentation we will provide a preview of the most relevant enhancements introduced.

New measurements in Db2 V13 (MT)

In this presentation we will discuss the most relevant new metrics available in Db2 13. We will also show which trace you have to activate to produce the needed IFCIDs and metrics, and the planned support in EPV zParser and EPV for Db2.

Subscription

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.

The subscription form is available at: www.epvtech.com

IBM Redbooks - Optimized Inferencing and Integration with AI on IBM Z

IBM Redbooks - Optimized Inferencing and Integration with AI on IBM Z
Introduction, Methodology, and Use Cases

This IBM Redpaper publication introduces and explains AI technologies on IBM Z and demonstrates the capabilities that are used in a use case that involves credit risk scoring on the platform.

It also describes and demonstrates the implementation and integration of an end-to-end solution, from developing and training a deep learning model, to deploying that model in an IBM z/OS V2R5 environment on an IBM z15 hardware, to integrating AI functions into an IBM z/OS CICS application.

Download it at: [Optimized Inferencing and Integration with AI on IBM Z](#)



Customer Questions

I see that EPV calculates two different capture ratios.
Can you please explain the capture ratio concept and how you calculate these two metrics?

EPV Technical Support answer

The system capture ratio is the ratio between the total CPU attributed to all the workload in a system and the total CPU used by the system itself.

The system capture ratio is an inverted measure of the system overhead. The higher the capture ratio the lower the system overhead.

A widely accepted ROT (Rule Of Thumb) for production systems is to consider as normal values between 85% and 95%.

The system overhead is the amount of CPU used for activities (such as paging, swapping, part of the I/O operation, etc.) not directly correlated to a specific address space, transaction or enclave.

The impact of these activities is not the same on all the workload types: a TSO is continuously swapped in and out while CICS and IMS regions are normally non swappable, some batch jobs do a huge amount of I/Os compared to their CPU usage while other applications work mostly in memory.

So, an important characteristic of the capture ratio is its workload dependence.

The two capture ratios provided by EPV for z/OS are calculated as follows:

- 30 capture ratio = $\text{CPU30} * 100 / \text{CPU70}$
- 72 capture ratio = $\text{CPU72} * 100 / \text{CPU70}$

where:

CPU70 is the total CPU used by the system; from SMF 70 records;

CPU72 is the total CPU used by all the WLM service classes in the system; from SMF 72 records;

CPU30 is the total CPU used by all the address spaces in the system; from the SMF 30 interval records (subtype 2 and 3).


In general, the two values are about the same, the 72 capture ratio being a bit higher because of a better accounting in SMF 72 records.

Please note that the 30 capture ratio can be misleading when there are address spaces swapped out for a long time. When this happens, their consumptions are accounted when they are swapped in. This could lead to intervals showing a lower capture ratio and intervals with a higher capture ratio compared to what is reported

by SMF 72 records.

In these cases, the 30 capture ratio could also be higher the 100 percent.

Little known SMF parameters



SIGSTRIP/NOSIGSTRIP

SIGSTRIP/NOSIGSTRIP is a parameter of both the IFASMF DL and IFASMF DP programs.
Default is SIGSTRIP.

It specifies whether or not the log stream dump program (IFASMF DL) or the SMF data set dump program (IFASMF DP) must include the signature records (SMF type 2 subtype 1 and 2) in the dump data sets.

With SIGSTRIP the signature records will not be included in the dump data sets (SMF type 2 records including the standard SMF header only will be written); with NOSIGSTRIP the signature records will be included in the dump data sets.

If no signature records are present in the input, there will be no impact on processing.

With IFASMF DP you must specify NOSIGSTRIP in order to use the SIGVALIDATE parameter.

Quotes



"Everything is hard before it is easy"

Goethe

Copyright © 2022 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 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 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 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 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
[why did I get this?](#) [unsubscribe from this list](#) [update subscription preferences](#)
EPV Technologies · Viale Angelico, 54 · Roma, RM 00195 · Italy

