



epv

IT Cost
Under Control

EPV Technologies

Newsletter

October 2021

THIS MONTH HIGHLIGHTS

- [Measuring the effects of System Recovery Boost - Part 2](#)
- [System Z Update virtual conference](#)
- [IBM z/OS 2.5 manuals](#)

Measuring the effects of System Recovery Boost Part 2

With the z15 machines, IBM provides an interesting new and free feature: the System Recovery Boost (SRB).

This feature is designed to diminish the impact of downtime, planned or unplanned, so you can restore service and recover workloads substantially faster than on previous IBM machine generations.

With SRB, for a fixed amount of time, you can get a performance boost which will allow you to reduce the length and mitigate the impact of downtime.

This performance boost is obtained by temporarily releasing the limitations imposed on “knee-capped” machines (sub capacity models) and on the amount of work which is marked as zIIP eligible.

A very important point to take note of is that SRB will not cause an increase in IBM software MSU consumption or cost.

In this paper we will discuss the most important SRB concepts and we will look at the new system parameters, commands and SMF fields introduced to control it.

We will then discuss the effects of SRB on CPU and zIIP measurements by showing a couple of real-life examples.

Finally, we will also present what we changed in EPV for z/OS to support this new feature.

If you want to receive the paper you can reply to this e-mail writing **"Measuring the effects of System Recovery Boost - Part 2"** in the subject

System Z Update virtual conference

System Z continues to evolve introducing new hardware and software technologies. In this virtual conference we will focus on some of the most recent and interesting of them.

The conference will be held, in cooperation with IBM, on November 24th and repeated on November 25th.

We are proud to announce that the ‘What’s new in z/OS 2.5’ session will be presented by Gary Puchkoff, current z/OS product manager.

In the recent years Gary has been creator of WebSphere Application Server on z/OS, inventor of z Application Assist Processor (fore runner of zIIP specialty engines), architect of z/OS Management Facility (zOSMF a browser UI for z/OS), and creator of z/OS Container Extensions (running Linux workloads directly in z/OS).

Agenda

Agenda			
24th and 25th November 2021			
08:55 – 09:00	Welcome and introduction	Danilo Gipponi	EPV Technologies
09:00 – 10:00	What's New in z/OS 2.5: GA Edition	Gary Puchkoff	IBM
10:30 – 11:00	Measuring the Integrated Accelerator for zEDC	Fabio Massimo Ottaviani	EPV Technologies
11:15 – 11:30	Coffee break		
11:30 – 12:00	IBM Z - GA2 latest news	Francesco Bertagnoli	IBM
12:30 – 13:00	Measuring zCX applications	Matteo Bottazzi	EPV Technologies
13:15	End of conference		

Subscription

The conference is free of charge and reserved to EPV customers, partners and invited guests.

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 z/OS 2.5 manuals

All the z/OS 2.5 manuals, in pdf format, are now available

Download it at: [z/OS V2R5](#)



We should like to know, if there is any way to know which Programs/Transactions are called from a specific WAS Server.
I suppose the SMF 120 records (subtype 9?) contains relevant information, but I am not that sure...

Anyway, thank you for any help!

EPV Technical Support answer

SMF 120 provides info about invoked objects (mostly servlets and EJBs).

In subtype 6 and 8 you have almost all the objects invoked but data are aggregated so for each object you know the total values (CPU and elapsed) and the number of executions in each interval. Intervals can be synchronized with SMF.

I say almost all the objects because as far as I know at least asynchronous beans information are not provided.

WebSphere creates one subtype 9 record for every request that the server processes — for both external requests (application requests) and internal requests.

You have detailed info only for the first 30 objects involved in the request; please note that to get object info you need to activate the CPU Usage Breakdown option.

In EPV for z/OS we use subtype 6 and 8.

For the EJB container (subtype 6) the application name is taken from the SM120JB1 field and the transaction name from the SM120JM1 field.

For the WEB container (subtype 8) the application name is taken from the SM120WAQ field and the transaction name from the SM120WIQ field.

For what concerns subtype 9, it is collected in the EPV zParser DB so you can get all the need info by running your queries.



Synchronous execution switch suspensions

Synchronous execution switch suspensions occur when the following operations have to be performed under a different execution unit than the one that is used by the thread:

- UPDATE COMMIT; it includes wait time for: log writes in the second phase of two-phase commit, log writes and database writes for LOB with LOG NO when doing single phase commit, Page P-locks unlocks for updated pages and GBP writes in data sharing environments
- OPEN/CLOSE; it includes wait time for dataset open/close;

- SYSLGRNG REC; it includes wait time for recording in SYSLGRNG;
- DATASET EXTEND/DELETE/DEFINE; it includes wait time for dataset define, extend, delete, reset and VSAM catalog access;
- OTHER SERVICE; it includes wait time for other services such as VSAM catalog update and DRDA requestor waiting for server response when using TCP/IP.

Most of the commit-related waits are caused by log writing.

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).

The logs should be allocated to the fastest disk subsystem available. In some cases, you can consider striping for these data sets.

A high number of open and close may indicate that your DSMAX threshold is frequently reached. You can check it by looking at the QTDSDRN field in SMF 101. If this is the case, you can consider changing CLOSE YES to CLOSE NO on data sets that are used by critical applications.

A record is inserted into SYSLGRNG on the first update to a data set after it is open. So, the number of open may also increase SYSLGRNG suspensions.

SYSLGRNG entries are also updated when a page set is converted from read-write state to read-only state. You can check it by looking at the QTDSDRN field in SMF 101.

If this is the case, you can consider increasing the PCLOSET and PCLOSEN parameters for data sets that are switching between read-only and read-write too frequently.

The most common reason for high DATASET EXTEND/DELETE/DEFINE suspensions is a slow access to the catalog.

You should check if there is a performance problem on the volumes and I/O paths where the catalog is. You should also check if your catalog definitions and settings are tuned for performance.

These considerations also apply to OTHER SERVICE suspensions in case of catalog updates.

Information about synchronous execution switch suspensions is provided at plan level (IFCID 003) in the following fields of SMF 101 records:


- QWACAWTE, accumulated wait time because of update commit;
- QWACARNS, number of wait trace events processed for update commit waits;
- QWAXOCSE, accumulated wait time because of dataset OPEN/CLOSE;
- QWAXOCNS, number of wait trace events processed for dataset OPEN/CLOSE;
- QWAXSLSE, accumulated wait time because of a switch to the DB2 SYSLGRNG recording service;
- QWAXSLNS, number of wait trace events processed for a switch to the DB2 SYSLGRNG recording service;
- QWAXDSSE, accumulated wait time because of DEFINE, EXTEND, DELETE, RESET dataset and VSAM catalog access;

- QWAXDSNS, number of wait trace events processed for DEFINE, EXTEND, DELETE, RESET dataset and VSAM catalog access;
- QWAXOTSE, accumulated wait time because of a switch to other Db2 services;
- WAXOTNS, number of wait trace events processed for a switch to other Db2 services.

Much less details are provided at package level (IFCID 239); you have just these two fields in SMF 101 records:

- QPACAWTE; accumulated wait time because synchronous execution switch; it includes UPDATE COMMIT, OPEN/CLOSE, SYSLGRNG REC, DATASET EXTEND/DELETE/DEFINE, OTHER SERVICE;
- QPACARNS, number of wait trace events processed for synchronous execution switch; it includes UPDATE COMMIT, OPEN/CLOSE, SYSLGRNG REC, DATASET EXTEND/DELETE/DEFINE, OTHER SERVICE.

Quotes



"The purpose of art is washing the dust of daily life off our souls"

Pablo Picasso

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 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

