

z/OS: Calculating Workload MSU Usage



John Burg
Carol Creter
Bradley Snyder
IBM

Version Date: March 13, 2020

This document can be found on the web, www.ibm.com/support/techdocs
Under the category of “White Papers.”

Table of Contents

Contents

| | |
|--|---|
| Table of Contents | 1 |
| Overview | 2 |
| Software MSU Calculation for Workload or Service Class | 3 |
| 1. Obtain APPL% For Service Class | 4 |
| 2. Obtain APPL% for Service Policy | 4 |
| 3. LPAR MSU from the CPU Activity Report | 5 |
| 4. Calculate the MSU | 5 |
| Summary | 5 |

Overview

There are different ways to measure the IBM Z System (IBM Z) processor: CPU seconds, percent utilization, MIPS, and service units. Service units have been used at many installations as a chargeback mechanism to have different business units pay for their fair share of the use of the processor.

Every IBM Z processor has defined SRM constants, one of these is the service units per second (SU/Sec). The SU/Sec rating is the number of service units a single CP can expect to deliver in one second. This constant is set to make z/OS installation specifications (such as IEAOPTxx parameters) transparent across the IBM Z processor range.

The service units used by a workload is simply the CPU time associated with that work multiplied by the SU/Sec value of the LPAR on which the work is running. For all workloads there are two CPU types, CPU time and SRB time. The total service units is the sum of both the CPU and SRB values.

CPU service = CPU x CPU service units;
SRB service = SRB x SRB service units.

When IBM started to price software based on capacity, MSU (million service units) was introduced as a metric to leverage the SU (service unit) designation for the processor that was always there. An MSU, million service units, is a measurement of the amount of processing work a processor can perform in one-hour, 3,600 seconds. The original MSU value would have been:

$(\text{Service unit per second} * 3600 \text{ seconds} * \# \text{ engines (GCP)}) / 1 \text{ million}$

With the introduction of the z990 in 2003 IBM also introduced the Mainframe Charter. The Mainframe Charter provided a framework for planned future investment and highlighted specific ways in which IBM intended to deliver ongoing value to customers.

With the z990, the mainframe charter introduced a program where IBM discounted the software which introduced concept of a processor having two different ratings - a software MSU and hardware MSU. Though the software MSU was discounted the hardware number remained consistent. This essentially broke the constant relationship between the MSU rating and the SU/Sec value of a processor. This discounting mechanism was continued for several processor generations so there is no constant ratio to go from SU/Sec to MSU used.

The SU/Sec and MSU values are generally used to charge for software that run on the IBM Z platform running on z/OS and by clients to charge lines of business for their processor consumption.

The issue here is there is only a single capacity designation for any processor but to due to variations in LPAR configuration, and workload characteristics the actual capacity of the processor varies.

In 2006 additional complexity was added with the introduction of specialty engines- zAAPs and zIIPS. The specialty engines can run IBM Z workload but do not change the model designation or MSU rating when added to IBM Z processor.

This leads to a challenge of being able to use service units (SU) which are being calculated using hardware MSU value and relate to a software MSU value. The challenge becomes acute if you need to determine the MSU used at a more granular level, such as how a workload or service class is driving MSU used. We need a simple method to convert the service units used by a service class or report class into MSU used at the LPAR level.

Software MSU Calculation for Workload or Service Class

The software MSU is not reported in the SMF records. To get the MSU used by a workload or a service class we need to create a ratio of CPU time for service class to CPU time for LPAR at the policy level multiply that by the total MSU used at the LPAR level.

The simplest method to do this is to use APPL% from the RMF Workload Activity Report as it does not contain zIIP time. The service or service time fields in the Workload Activity Report will have a combined total of CPU, SRB, and zIIP time. The APPL% field only contains CPU and SRB time.

The APPL% is Percentage of the processor time used by transactions running on general purpose processors in the service or report class period. The RMF Workload Activity Report will have the APPL% for the service class as well as at the policy level, which is the total CPU used for the partition.

The MSU used for the partition can be retrieved from the CPU Activity Report. The final calculation is as follows:

$$\text{MSU(Workload)} = \text{MSU(LPAR)} * \frac{\text{APPL\%(Workload)}}{\text{APPL\%(Policy)}}$$

The following is an example of the steps needed to get the fields for the above calculation. We are trying to calculate the MSU used by service class BATCLASS on LPAR IBM999 (SYSID IBM9).

1. Obtain APPL% For Service Class

The following is a section of the Workload Activity Report run at the service class level. The APPL% for this service class is in the middle column of the second section as shown below.

The APPL% for this service class is 74.42

| POLICY=WLMBASE | | | WORKLOAD=BATCH | | | SERVICE CLASS=BATCLASS | | | |
|-----------------|--------|------------|------------------------------|-------|-------|------------------------------------|-------|-------------|-------|
| | | | | | | CRITICAL =NONE | | | |
| -TRANSACTIONS-- | | | TRANS-TIME HHH.MM.SS.FFFFFFF | | | TRANS-APPL%-----CP-IIPCP/AAPCP-IIP | | | |
| AVG | 1.07 | ACTUAL | 17.976323 | | | TOTAL | 73.40 | 0.00 | |
| MPL | 1.07 | EXECUTION | 17.357236 | | | MOBILE | 0.00 | 0.00 | |
| ENDED | 55 | QUEUED | 619087 | | | CATEGORYA | 0.00 | 0.00 | |
| END/S | 0.06 | R/S AFFIN | 0 | | | CATEGORYB | 0.00 | 0.00 | |
| #SWAPS | 2 | INELIGIBLE | 0 | | | | | | |
| EXCTD | 0 | CONVERSION | 57195 | | | | | | |
| | | | STD DEV | | | 14.343292 | | | |
| ----SERVICE---- | | | SERVICE TIME | | | ---APPL %--- | | | |
| IOC | 526229 | CPU | 578.774 | CP | 74.42 | BLK | 0.000 | SSCHRT 2846 | AVG |
| CPU | 20353K | SRB | 81.789 | IIPCP | 0.00 | ENQ | 0.000 | RESP 0.0 | TOTAL |
| MSO | 5083K | RCT | 0.002 | IIP | 0.00 | CRM | 0.000 | CONN 0.0 | SHARE |
| SRB | 2876K | IIT | 9.185 | AAPCP | 0.00 | LCK | 6.398 | DISC 0.0 | |
| TOT | 28838K | HST | 0.002 | AAP | N/A | SUP | 0.000 | Q+PEND 0.0 | |
| /SEC | 32042 | IIP | 0.000 | | | | | IOSQ 0.0 | |
| ABSRPTN | 30K | AAP | N/A | | | | | | |
| TRX SERV | 30K | | | | | | | | |

2. Obtain APPL% for Service Policy

The following is a section of the Workload Activity Report run at the policy level. The APPL% at this level will be the total CPU used for all workloads on the partition.

The APPL% for the partition at the policy level is 141.00.

| POLICY=WLMBASE | | | | | | | | | |
|-----------------|--------|------------------------------|----------|--------------|------------------------------------|--------|---------------|--------|-------|
| -TRANSACTIONS-- | | TRANS-TIME HHH.MM.SS.FFFFFFF | | | TRANS-APPL%-----CP-IIPCP/AAPCP-IIP | | | | |
| AVG | 175.72 | ACTUAL | | 550284 | TOTAL | 139.27 | 0.01 | | |
| MPL | 175.72 | EXECUTION | | 46533 | MOBILE | 0.00 | 0.00 | | |
| ENDED | 21805 | QUEUED | | 2694 | CATEGORYA | 0.00 | 0.00 | | |
| END/S | 24.23 | R/S AFFIN | | 0 | CATEGORYB | 0.00 | 0.00 | | |
| #SWAPS | 4940 | INELIGIBLE | | 0 | | | | | |
| EXCTD | 0 | CONVERSION | | 336 | | | | | |
| | | STD DEV | | 28.229469 | | | | | |
| | | | | | | | | | |
| ----SERVICE---- | | SERVICE TIME | | ---APPL %--- | --PROMOTED-- | | --DASD I/O--- | | ----S |
| IOC | 603420 | CPU | 1111.378 | CP | 141.00 | BLK | 0.000 | SSCHRT | 3488 |
| CPU | 39081K | SRB | 169.094 | IIPCP | 0.01 | ENQ | 0.000 | RESP | 0.1 |
| MSO | 8255K | RCT | 0.372 | IIP | 1.34 | CRM | 0.000 | CONN | 0.0 |
| SRB | 5946K | IIT | 15.214 | AAPCP | 0.00 | LCK | 9.252 | DISC | 0.0 |
| TOT | 53886K | HST | 0.010 | AAP | N/A | SUP | 0.000 | Q+PEND | 0.0 |
| /SEC | 59873 | IIP | 12.063 | | | | | IOSQ | 0.0 |
| ABSRPTN | 341 | AAP | N/A | | | | | | |
| TRX SERV | 341 | | | | | | | | |

3. LPAR MSU from the CPU Activity Report

The following is a section of the CPU Activity Report. The MSU used by the partition is highlighted below.

The MSU used by this partition is 117.

```

z/OS V2R3                      SYSTEM ID IBM9          START 07/01/2019
                                RPT VERSION V2R3 RMF      END   07/01/2019

MVS PARTITION NAME             IBM999          PHYS PROC NUM    14          G
IMAGE CAPACITY                  310                      CP           8          L
NUMBER OF CONFIGURED PARTITIONS 8                      IFL          2          A
WAIT COMPLETION                  NO                      ICF          2
DISPATCH INTERVAL              DYNAMIC                  IIP          2

----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA
-----MSU----- --CAPPING--- --PROCESSOR ----DISPATCH TIME DATA-
NAME      S  WGT  DEF    ACT  DEF    WLM%   NUM  TYPE  EFFECTIVE    TOTAL
IBM999    A   30    0    117  N N N    0.0    4  CP    00.22.26.372 00.22.36.
IBM444    A   15    0    29  N N N    0.0    3  CP    00.05.30.258 00.05.33.
IBM222    A   50    0   111  N N N    0.0    6  CP    00.21.15.687 00.21.26.
IBM777    A    5    0    11  N N N    0.0    3  CP    00.02.00.427 00.02.02.
*PHYSICAL*
-----

```

4. Calculate the MSU

Remember the original formula:

$$\text{MSU(Workload)} = \text{MSU(LPAR)} * \frac{\text{APPL\%(Workload)}}{\text{APPL\%(Policy)}}$$

From this we have all the values we need.

$$\text{MSU} = 117 * (74.42/141) = 61.75$$

In this example, of the 117 MSUs used by the partition, the BATCLASS service class used 61.75 MSUs.

Summary

If you need to calculate an MSU for a specific workload there is a simple calculation to get this done. You can use the ratio of CPU time for the workload to the total CPU time for the policy and multiplying this by the MSU used at the LPAR level.