

# **z/OS: Calculating the Capture Ratio**



John Burg  
Bradley Snyder  
IBM

Version Date: June 29, 2017

This document can be found on the web, [www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)  
Under the category of “White Papers.”

# Table of Contents

## Contents

Table of Contents .....	1
Overview .....	2
Overview - Capture Ratio Definition.....	3
RMF Report Examples .....	3
CPU Activity Reports .....	3
Workload Activity Report.....	5
Capture Ratio Calculations .....	7
Traditional Method – No Specialty CPs .....	7
1. GCPs Only – Using Service Time:.....	7
2. GCPs Only – Using APPL%: .....	7
Processor with Specialty CPs.....	7
1. Full Speed GCPs with zIIPs – Using Service Times:.....	8
2. Sub-capacity GCPs with zIIPs – Using Service Times: .....	8
3. GCPs and zIIPs – Using APPL %: .....	8
Processor with zIIPs and SMT .....	9
1. SMT – Using Service Time: .....	9
2. SMT and sub-capacity Processor – Using Service Time: .....	9
3. SMT – Using APPL%: .....	9
Observations on Calculation with SMT Enabled.....	9
Summary .....	10

## Overview

The system capture ratio in a z/OS system is a key performance metric which can help identify properly running systems. This key measure should always be one of the first data points examined when performing any system performance check using RMF data.

Before the introduction of specialty CPs as well as Simultaneous Multithreading (SMT) on modern z/OS processors calculating the capture ratio was a rather simple process. With the introduction of these new technologies, the amount of work running on a general purpose CP (GCP) may change based on specialty CP capacity, and workload arrival rates. This could cause the observed capture ratio to change even though there are really no changes occurring on the system.

Today a simple single capture ratio needs to consider the effects of these additions in order to come up with one all-inclusive number for the system.

This paper will go over the ways to calculate the capture ratio for each kind of configuration. It will start with the basic method which has been in place and go through the different configuration options and how each one affects the calculation of the system capture ratio.

# Overview - Capture Ratio Definition

There will always be some time spent in the operating system in which the system is doing work on behalf of other work requests, but this time cannot be directly attributed to any specific workload, transaction, or job. This system work is, in essence, 'uncaptured' and is the difference between a partition's application work CPU time and a partition's CPU time. Simply put, the capture ratio compares the amount of CPU time attributed to all measurable work in a partition with the actual amount of CPU time the partitions was dispatched to physical processors. This can be shown in the simple formula below

$$\text{CAPTURE RATIO} = \text{WorkCPU} / \text{PartitionCPUTime}$$

In the examples below this paper will use examples and data from IBM's Resource Measurement Facility (RMF) to show all the calculations necessary for calculating the capture ratio. In all examples below, all specialty CP time will be shown as zIIP time.

All calculation below will result in a number less than 1. To get the capture ratio in percent format multiple any of the results by 100%.

## RMF Report Examples

### CPU Activity Reports

The following CPU Activity report snapshot highlights three main data points of interest in this report. When a processor is running with SMT enabled for zIIP processors the capture ratio will need to take multi-threading time into account. The fields used for these are the Maximum Capacity Factor (MAX CF), Capacity Factor (CF), and the MT PROD values for zIIPs in the report below. In the report below the values for MAX CF are 1.490, CF of 1.281 and MT PROD of 85.92.

```

C P U   A C T I V I T Y

      z/OS V2R1          SYSTEM ID SYSD      DATE 03/19/2015      INTERVAL 04.59.999
RPT VERSION V2R1 RMF          TIME 13.40.00          CYCLE 1.000 SECONDS

-CPU---
NUM TYPE  ONLINE  LPAR BUSY  MVS BUSY  PARKED  --- MT % ---  LOG PROC  --
A  IIP    100.00  87.30     74.56     0.00    87.60  76.48  100.0  HIGH
      68.51     0.00
B  IIP    100.00  81.12     67.77     0.00    85.35  69.24  100.0  HIGH
      58.38     0.00
C  IIP    100.00  57.47     47.81     0.00    84.57  48.61  53.2   MED
      39.82     0.00
D  IIP    100.00  29.29     25.03     0.00    84.67  24.80  53.2   MED
      20.02     0.00
E  IIP    100.00   9.15      8.00     0.00    87.40   7.99   0.0   LOW
      6.66      0.00
TOTAL/AVERAGE          52.87      41.66          85.92  45.42  306.4
0----- MULTI-THREADING ANALYSIS -----
      CPU TYPE  MODE  MAX CF      CF      AVG TD
      CP        1      1.000      1.000      1.000
      IIP       2      1.490      1.281      1.573
  
```

The next section of the CPU Activity Report is the Partition Data report. This will be where an LAPR's physical dispatch time and percentage of CEC utilization by CP type are located.

The primary data points obtained from the partition data report are the Dispatch Times for the GCPs and the Dispatch Time for the zIIPs. In the example below the partition UOSP02's GCP Dispatch Time is 37.084 seconds and the partition UOSP02's zIIP Dispatch Time is 793.009 seconds (converted from 13 minutes and 13.009 seconds in the report).

In some cases we will also use the Physical Processors Total Utilization Percentage for GCPs and zIIPs as GCP% and zIIP% respectively. In this case these values are 0.34 for GCP% and 22.03 for zIIP%

```

1          P A R T I T I O N   D A T A   R E P O R T
z/OS V2R1  SYSTEM ID SYSD      DATE 03/19/2015      INTERVAL 04.59.999
...
-- PARTITION DATA-  -- LOGICAL PARTITION PR--  -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
      -PROCESSOR  ----DISPATCH TIME DATA----  LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
      NAME  NUM  TYPE  EFFECTIVE  TOTAL  EFFECTIVE  TOTAL  LPAR MGMT  EFFECTIVE  TOTAL
UOSP02    6  CP    00.00.33.360  00.00.37.084    1.85    2.06    0.03    0.31  0.34
.....
UOSP02    5  IIP  00.13.09.585  00.13.13.009   52.64   52.87    0.10   21.93  22.03
.....

```

## Workload Activity Report

The following is a Workload Activity Report snapshot. When calculating the capture ratio we want the total CPU time for all workloads running on the system. Therefore, from the workload activity report we want the data from the policy level of the report, as shown below. The two columns where we are getting data are the SERVICE TIME and the APPL % columns.

To get the Total CPU service time from the SERVICE TIME column (Total CPUST) the sum of CPU, SRB, RCT, IIT, and HST will need to be summed.

$$\text{Total CPUST} = \text{CPU} + \text{SRB} + \text{RCT} + \text{IIT} + \text{HST}$$

The zIIP time (IIP) is already included in the CPU field and will not have to be into the Total CPU service time. If the partition is running on a sub-capacity model processor or the zIIPs are running in SMT mode, the IIP time has been normalized in either the CPU or the IIP fields. This normalization is to take into account either the speed differences between processor types, or converting to MT=1 equivalent time, or both. The calculations below will show how to handle these special cases

The CP APPL% and IIP APPL % fields include all recorded time spent on each processor type. On a sub-capacity model processor no adjustments are necessary when using these fields, but with SMT on the zIIPs some adjustment will be necessary. This will be shown in detail later.

In the report below, the Total CPUST = 1,031.837. This is the sum of CPU+SRB+RCT+IIT+HST. The CPU time includes the IIP service time (1009.730) which is why it does not need to be included in the Total CPUST calculation.

The CP APPL% is 7.37 and the IIP APPL% is 225.86. These fields, with the CPU fields mentioned above, would be used in the calculations below for this system's capture ratio.

```

1                               W O R K L O A D   A C T I V I T Y

z/OS V2R1      SYSPLEX WSCZPLEX   DATE 03/19/2015   INTERVAL 04.59.999   MODE = GOAL
                RPT VERSION V2R1 RMF                TIME 13.40.00

----- SERVICE POLICY

REPORT BY: POLICY=WLMPOL
                WSC WLM test policy

-TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  ---SERVICE---  SERVICE TIME  ---APPL %---
AVG           75.57  ACTUAL                17  IOC           703139  CPU 1027.474  CP           7.37
MPL           75.57  EXECUTION                17  CPU           76463K  SRB   3.965  AAPCP       0.00
ENDED        74121  QUEUED                   0  MSO              0  RCT   0.002  IIPCP       1.97
END/S        247.07  R/S AFFIN                0  SRB           295095  IIT   0.406
#SWAPS         38  INELIGIBLE              0  TOT           77461K  HST   0.000  AAP         N/A
EXCTD         0  CONVERSION              0  /SEC          258205  AAP         N/A  IIP       225.86
AVG ENC       5.25  STD DEV                  14  IIP 1009.730
REM ENC       0.00
MS ENC       0.00
                ABSRPTN 3417
                TRX SERV 3417

```

In some environments work is running on what is called a sub-capacity processor. In these cases the specialty CPs have a different capacity setting while the specialty CPs are still running at full speed. In these cases there is a normalization factor (zNF) which will adjust specialty CP times to an equivalent GCP time for reporting purposes. When running on a sub-capacity processor, some of the calculations will need to take this normalization factor into account when calculating the combined system capture ratio.

The normalization factor value is kept in the RMF 72 record and is reported in the Workload Activity Report on the Service Policy Page, as seen below. In this example the zNF is 1.0000, meaning the processor is not a sub-capacity model.

```

- SERVICE POLICY PAGE -

SERVICE DEFINITION: WLM          WSC  -SERVICE DEFINITION COEFFICIENTS-
INSTALL DATE: ... MCCOX          IOC   CPU   SRB   MSO  -NORM FACTORS-
POLICY: WLMPOL          WSC policy
I/O PRIORITY MANAGEMENT: YES      0.1   1.0   1.0  0.0000  AAP   IIP
DYNAMIC ALIAS MANAGEMENT: YES
1.0000  1.0000

```

## Capture Ratio Calculations

The examples below will show the calculation using both CPU service times and utilization based on core or CP utilizations. Either method works and will provide equivalent answers so the choice of method is left to the user.

### ***Traditional Method – No Specialty CPs***

In a simple environment there will be no specialty CPs (zIIPs or zAAPs). In this case the capture ratio is simple the ratio of the total of the CPU time in the RMF 72 records in all service classes to the CPU time in the RMF 70 records for the partitions.

#### **1. GCPs Only – Using Service Time:**

$(\text{Total CPUST})/(\text{Partition GCP Time})$

Where CPUST is the total of all CPU Service times reported in the RMF Workload Activity Report at the policy level and Partition GCP Time is the total dispatch time for the partition as reported in the RMF CPU Activity Report.

#### **2. GCPs Only – Using APPL%:**

$(\text{CP APPL\%})/(\text{Number GCPs} * \text{GCP \%})$

We can also use the CP APPL% field from the RMF Workload Activity Report with the physical busy percentage for the partition from the CPU Activity Report. There will be a need to adjust the physical busy by the number of CPs on the processor in order to have both measurements on the same scale. The APPL% value is based on percent of a single CP whereas physical CP busy is based on the busy of the entire CEC. The simplest method is to multiple the Physical Busy with the number of GCPs on the processors

### ***Processor with Specialty CPs***

When specialty CPs are added to the system there will be two separate CPU buckets for the system. The first will be all CPU time for the GCPs and the second will be the CPU time associated with the specialty CPs. In our examples all specialty CP times will be zIIP processors.

There are two possibilities for how the processor is running with zIIPs. The first case is all GCPs and specialty CPs are running at the same speed. In this case the processor model will be a 7xx model. The second case is the GCPs are running as a sub-capacity model, in which case the specialty CPs may be running at a much faster speed. The sub-capacity processor models include 4xx, 5xx, and 6xx models. The CPU Service Time in the RMF Workload Activity Report for specialty CPs and in the RMF 72 records is adjusted to an equivalent GCP time in order to have a consistent base for the measurement.

When the GCP processors and specialty processors are running at the same speed, the formula is a simple addition of the different dispatch times from the report. The formula below assumes all specialty CP time is zIIP time. If there are other specialty CPs, such as zAAPs then that time would be added as well.

### **1. Full Speed GCPs with zIIPs – Using Service Times:**

$(\text{Total CPUST}) / (\text{Partition GCP Time} + \text{Partition zIIP Time})$

When using the service times from the RMF Workload Activity Report, all rows in the service time column will need to be added to CPU except the row IIP service time as it is already included in the total CPU Service time.

When using the CPU time from the RMF CPU Activity Report the Partition GCP Time and Partition zIIP Time for each processor type need to be added together.

### **2. Sub-capacity GCPs with zIIPs – Using Service Times:**

$(\text{Total CPUST}) / (\text{Partition GCP Time} + (\text{Partition zIIP Time} * \text{zNF}))$

When the GCPs are running as a sub-capacity model, then the formula for the capture ratio will need to take the difference in uni-processor speed into account. There will be more time included in the reported CPU time than was dispatched because of the adjustments made by the zIIP normalization factor (zNF). In this case we will adjust the zIIP partition dispatch time to match the normalization that was done in the service class level.

### **3. GCPs and zIIPs – Using APPL %:**

$(\text{CP APPL\%} + \text{zIIP APPL\%}) / ((\text{Number GCPs} * \text{GCP \%}) + (\text{Number zIIPS} * \text{zIIP \%}))$

Lastly, we can get the total capture ratio by using the APPL% values for each processor type from the workload activity report and the total physical utilization from the CPU Activity report. Again, the physical utilization for each CP type will need to be multiplied by the number of CPs of that type to get the same scale as the APPL% values.

A benefit of using the utilization percentages from the RMF report means no adjustments are necessary when running on a sub-capacity model processor. The zIIP APPL% value is the actual percentage of a single CP used by the system.

## **Processor with zIIPs and SMT**

With SMT, two threads can execute at the same time on a zIIP processor. There is a new calculation, called the 'Capacity Factor' (CF) which measures the additional capacity gained with SMT for the zIIP processors. With SMT enabled the amount of CPU time accumulated on zIIPs will need to be adjusted by this capacity factor in order to normalize this time back to an equivalent GCP time.

### **1. SMT – Using Service Time:**

$(\text{Total CPUST}) / (\text{Partition GCP Time} + (\text{Partition zIIP Time} * \text{CF}))$

In this case, the numerator of the formula is the same as when running with zIIPs. The difference is the reported zIIP time from the 70 record, or the CPU Activity report, must be adjusted by the Capacity Factor. This adjustment is necessary as the zIIP service times in the 72 records and the Workload Activity Report have been adjusted and therefore both times need to be on the same scale.

### **2. SMT and sub-capacity Processor – Using Service Time:**

$(\text{Total CPUST}) / (\text{Partition GCP Time} + (\text{Partition zIIP Time} * \text{zNF} * \text{CF}))$

When the GCPs are running as a sub-capacity model, the capacity . The formulas for full speed and sub-capacity processors are really the same. The only difference is when the processor is a full capacity model, the zIIP normalization factor is always 1.

### **3. SMT – Using APPL%:**

$(\text{CP APPL\%} + \text{zIIP APPL\%}) / ((\text{Number GCPs} * \text{GCP \%}) + (\text{Number zIIPs} * \text{zIIP\%} * \text{MT PROD\%}))$

We can still use the CP APPL% and zIIP APPL% values in any capacity model type. In this case the formula takes the capacity factor into account by using the MT PROD value from the RMF report and applying it to the zIIP time from the CPU Activity report. The MT Prod value is simply the ratio of the CF to the Max CF.

There is not much change from this formula to the same formula when there is no SMT enabled. Without SMT enabled the MT PROD value is always 100%.

## **Observations on Calculation with SMT Enabled**

With zIIPs exploiting SMT, at very low and high utilizations MT metrics, and therefore the capture ratio calculations, may produce unexpected results. It is recommended to

calculate both a system capture ratio as well as a GCP only capture ratio using the APPL% value calculation when running with SMT enabled.

## Summary

By using the examples above, a capture ratio for any configuration can be calculated in order to give a correct view for the overall system. When using RMF reports the simplest method will be using the percent utilization and APPL% values for each processor type as there is no change needed when running on sub-capacity models or when running with or without SMT enabled (MT PROD = 100% without SMT). The table below shows the formulas to use for each scenario.

Without SMT	$(\text{CP APPL\%} + \text{zIIP APPL\%}) / ((\text{Number GCPs} * \text{GCP \%}) + (\text{Number zIIPs} * \text{zIIP\%}))$
With SMT	$(\text{CP APPL\%} + \text{zIIP APPL\%}) / ((\text{Number GCPs} * \text{GCP \%}) + (\text{Number zIIPs} * \text{zIIP\%} * \text{MT PROD\%}))$

Both formulas are the same. The only difference is when SMT is not enabled, MT PROD is always 100%. When no specialty CPs are present, the number of zIIPs and zIIP APPL% will be zero.