



## WLM: the SYSSTC service class

*Fabio Massimo Ottaviani – EPV Technologies*

November 2014

### 1 Introduction

The Workload Manager (WLM) is an essential component of the z/OS environment. Its role is to classify the system workloads in user defined service classes and to manage those workloads towards the importance and the goal set for each service class period by the user.

However there are two service classes which are not created by the user: SYSTEM and SYSSTC<sup>1</sup>. These service classes always exist in a z/OS system.

An important characteristic of these service classes is that they don't have a goal so you can't measure a Performance Index for them. They are managed with a very high fixed Dispatching Priority (DP) and I/O Priority (IOP):

- SYSTEM gets 255 for both DP and IOP; this is the highest possible priority; it is used for operative system address spaces;
- SYSSTC gets 254 for both DP and IOP; it is generally used for highly important started tasks processing.

In this paper we will focus on the SYSSTC service class to discuss:

- what is classified in SYSSTC by default;
- what the user can and should do;
- what performance measurements are available and how they can be used.

---

<sup>1</sup> A third system-provided service class SYSOTHER is used as default service class for non-STC address spaces when no classification rules exist for the subsystem type. It is assigned a discretionary goal.



## 2 What is classified in SYSSTC by default

Only started tasks (STC) running a program with the PRIV (privileged) or SYST (system task) special attributes will be assigned to the SYSSTC service class by default.

By customizing the PPT (Program Property Table) parameter inside the SCHEDxx member of the SYS1.PARMLIB library the installation may specify special attributes for specific programs or override an IBM-supplied entry in the PPT<sup>2</sup>.

Customers should not add or change program attributes in the program properties table (PPT) unless specifically required when installing a program. Also in this case such an action should be carefully evaluated before performing it.

More details can be found in the “MVS Initialization and Tuning Reference” manual.

Starting from z/OS 2.1 a Display PPT command is provided to list default and active PPT.

An example of the command output, taken from the “MVS System Commands” manual, is shown in Figure 1. Only part of the rows is presented.

```

IEF386I 11.20.01 DISPLAY PPT
No Parmlib Values
Default Values
PgmName  NC NS PR ST ND BP Key 2P 1P NP NH CP
AHLGTF   Y Y . Y . . 0 . . Y . .
AKPCSIEP . Y . Y Y . 1 . . Y . .
ANFFIEP  . Y . Y Y . 1 . . . . .
APSHPOSE . Y . Y Y . 1 . . Y . .
APSKAFPD . Y . Y Y . 1 . . Y . .
APSPPIEP . Y . Y Y . 1 . . Y . .
ASBSCHIN . Y . Y . . 1 Y Y . . .
ASBSCHWL . . Y . . . 1 . . . . .
ATBINITM . Y . Y . . 1 Y Y . . .
ATBSDFMU . . Y . . . 1 . . . . .
AVFMNBLD Y Y . Y . . 3 . . Y . .
BBGCTL   . Y Y . . . 2 . . . . .
BBGDAEMN . Y Y Y . . 2 . . . . .
BBOCTL   . Y Y . . . 2 . . . . .
BBODAEMN . Y Y Y . . 2 . . . . .
BNJLINTX . Y . . . . 8 . . . . .
    
```

Figure 1

## 3 What the user can and should do

The SYSSTC service class is designed to be used for workloads that use a small amount of resources (mostly CPU) but need to run with the smallest possible delays. A perfect example is the DB2 IRLM address space.

To avoid delays, SYSSTC workloads run at a fixed Dispatching Priority set to 254 which is the second highest priority in the system.

So tasks in SYSSTC should not use a lot of CPU to avoid monopolizing the CPU resources with consequent delays for all the other workloads.

<sup>2</sup> IBM supplied entries are provided in the IEFSDPPT member of the SYS1.LINKLIB library.



To minimize the risk of loops we also suggest to put in SYSSTC only workloads running “trusted code”; an example of “trusted code” is the IMS control region while an example of “non-trusted code” is an IMS message region running user applications<sup>3</sup>.

The idea behind is that, while it’s always possible a bug, IBM and ISV code is normally more tested in many different environments than the specific user applications.

We also suggest putting in SYSSTC workloads not using large amounts of memory, unless memory is really very abundant in the system; the reason is that WLM doesn’t protect the working set of address spaces running in SYSSTC (and SYSTEM)<sup>4</sup>. So if memory is constrained a lot of memory frames belonging to SYSSTC address spaces may be paged-out and then paged-in when needed hurting the overall system performance.

Examples of address spaces normally running with a very large memory footprint are DB2 DBM1, WebSphere MQ MSTR, WebSphere Servant regions.

In contrast with what we discussed about the SYSTEM service class (see EPV Newsletter 10/2014) there are some actions a customer should do in WLM definitions in order to change the default classification:

- setting STC classification rules to be sure that address spaces using a lot of resources (CPU and memory) don’t run in SYSSTC;
- setting STC classification rules to be sure that the appropriate address spaces run in SYSSTC;
- assigning a specific report class to the system address spaces you want to keep under control;
- setting an SPM rule in STC classification rules to be sure that address spaces not classified before and running programs with the PRIV and SYST attributes run in SYSSTC.

```

Subsystem-Type  Xref  Notes  Options  Help
-----
Command ==>>  Modify Rules for the Subsystem Type  Row 54 to 60 of 60
                Scroll ==>> CSR

Subsystem Type . : STC          Fold qualifier names?  Y (Y or N)
Description . . . STARTED TASK

Action codes:   A=After      C=Copy      M=Move      I=Insert rule
                B=Before      D=Delete row R=Repeat    IS=Insert Sub-rule
                More ==>>

-----Qualifier-----
Action  Type  Name  Start  Service  Report
-----
DEFAULTS: STCLOW  RSTC
_____ 1 TN   DB2TMSTR  _____  IMSDB2LO  DB2TMSTR
_____ 1 TN   DB2TDBM1  _____  IMSDB2LO  DB2TDBM1
_____ 1 TN   DB2TIRLM  _____  SYSSTC    DB2TIRLM
_____ 1 TN   DB2PMSTR  _____  IMSDB2HI  DB2PMSTR
_____ 1 TN   DB2PDBM1  _____  IMSDB2HI  DB2PDBM1
_____ 1 TN   DB2PIRLM  _____  SYSSTC    DB2PIRLM
_____ 1 SPM  SYSSTC    _____  SYSSTC    RSYSSTC
***** BOTTOM OF DATA *****
    
```

Figure 2

<sup>3</sup> A similar example is a CICS TOR (trusted) switching all transactions to some CICS AORs (not trusted).

<sup>4</sup> WLM protects only the working set of address spaces running in service classes with user importance (from 1 to 5) and goal.



In the example in Figure 2:

- we set STC classification rules to be sure that DB2 MSTR and DBM1 address spaces don't run in SYSSTC;
- we set STC classification rules to be sure that DB2 IRLM address spaces run in SYSSTC;
- we assigned specific report classes to each DB2 address space to keep it under control;
- we set an SPM rule in STC classification rules at the very end to be sure that address spaces, not classified before and running programs with the PRIV and SYST attributes, run in SYSSTC.

The SYSSTC service class can be assigned also to work running in the ASCH, JES, OMVS and TSO subsystems; so you should check if you need to set similar rules also for them.

It's enough common assigning one or more TSO users to SYSSTC in order to be able to quickly execute commands when the system is slowing down.

An example of this TSO classification rule is provided in Figure 3.

Subsystem-Type Xref Notes Options Help						
Modify Rules for the Subsystem Type						Row 1 to 8 of 33
Command ==>						Scroll ==> CSR
Subsystem Type . . :		TSO		Fold qualifier names?		Y (Y or N)
Description . . .		TSO				
Action codes: A=After C=Copy M=Move I=Insert rule						
B=Before D=Delete row R=Repeat IS=Insert Sub-rule						
						More ==>
-----Qualifier-----			-----Class-----			
Action	Type	Name	Start	Service	Report	
				DEFAULTS: ALLTSO TSO		
---	1	UIG		SYSSTC		
---	1	AI	349*	2		TSO349
---	1	AI	802*	2		TSO802
---	1	AI	805*	2		TSO805
---	1	AI	807*	2		TSO807
---	1	AI	810*	2		TSO810
---	1	AI	812*	2		TSO812
---	1	AI	813*	2		TSO813

Figure 3

Another common issue is about monitor address spaces. Do you have to put them in SYSSTC ?

The answer is: it depends.

RMF address spaces can normally be considered “trustable”. Other monitors such as Mainview and Omegamon could be a bit more dangerous depending on how you exploit them (parameters, traces, etc.).

Since many years the service classes SYSSTC1, SYSSTC2, SYSSTC3, SYSSTC4, and SYSSTC5 are provided; IBM manuals say “for future use”.

Work assigned to any of these service classes is managed identically to work assigned to SYSSTC. Currently, there is no technical reason to choose one of these service classes as an alternative to SYSSTC.



## 4 What performance measurements are available and how they can be used

As already mentioned, the SYSSTC service class doesn't have a goal to reach. So it's not possible to analyze its performance by looking at the Performance Index (PI) as you do normally for user defined service classes.

You can get an idea about the SYSSTC performance by looking at its velocity.

From the point of view of resource usage you have many metrics available in SMF 72 about CPU, zIIP, memory, I/O, etc.

Similar considerations as the one discussed for SYSTEM velocity and CPU usage can be applied to SYSSTC. The major difference is that workload running in SYSSTC is generally less stable and, also depending on customer's classification rules, uses more CPU than workload in SYSTEM.

Please refer to the EPV Newsletter 10/2014 for details and examples.

In the next chapter we will discuss SYSSTC memory usage.

### 4.1 SYSSTC memory usage

In the figure below we reported a screenshot of an EPV for z/OS view where we selected only the first 12 hours of the day. The table is sorted by descending values at hour 7<sup>5</sup>.

You can see that the TOTAL memory used by all the service classes rose from 16,3 to almost 18 GB. The major memory user is the WAS5AS service class and the second is SYSSTC.

SRVCLASS	PERIOD	0	1	2	3	4	5	6	7 ↓	8	9	10	11
TOTAL		16.413,2	16.307,5	16.285,5	16.284,4	16.291,0	16.333,6	16.361,3	16.852,6	17.749,4	17.546,7	17.177,4	17.098,3
WAS5AS	1	7.264,9	6.894,5	6.828,8	6.747,5	6.715,4	6.452,2	6.434,0	8.535,2	8.626,5	8.304,5	7.323,9	7.014,5
SYSSTC	1	5.116,9	5.002,1	4.960,4	4.905,6	4.927,6	4.682,7	4.454,3	4.829,1	5.004,2	4.722,5	4.423,4	4.475,9
IMSDB2LO	1	2.634,4	2.835,0	2.888,9	2.895,5	2.893,2	2.866,3	2.824,0	1.875,1	2.571,7	2.703,4	3.067,7	3.213,1
STCLOW	1	905,9	840,1	879,1	981,8	1.025,6	1.606,4	1.924,9	1.102,9	1.089,3	1.336,3	1.843,5	1.844,6
SYSTEM	1	396,6	575,2	587,2	578,4	581,7	578,4	575,9	421,6	383,0	401,3	434,1	454,8

Figure 4

At 7, WAS5AS used memory increased by about 2GB while SYSSTC shows an increase of less than 400 MB. All the other service classes show a reduced memory footprint.

Unfortunately there is no information at the address space level in SMF.

To understand how much memory each address space uses the simplest possibilities are:

- set a specific report class for each SYSSTC address space you want to control;
- look at RMF Monitor III STORF report; by setting the report options for STORF you can get only address spaces running in SYSSTC.

<sup>5</sup> A similar view is provided for report classes.



If you use the first option you only need to look at the REPORT CLASS MEMORY USAGE view in EPV for z/OS.

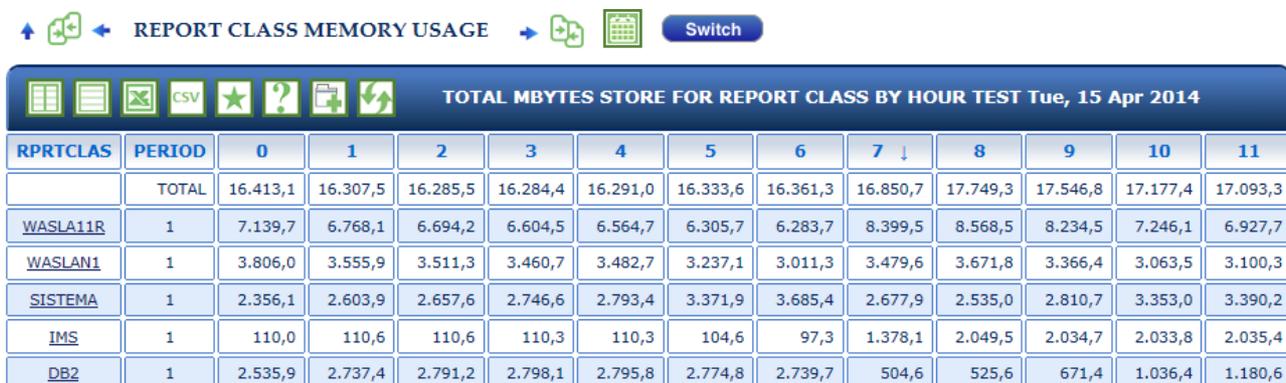


Figure 5

The second address space (WASLAN1) in Figure 5 runs in SYSSTC. It is a WebSphere address spaces. You can note it normally uses more than 3 GB and it increased its memory footprint by about 400 MB at 7.

## 5 Conclusions

Default assignment of SYSSTC address spaces is not always optimal. In this paper we discussed the most relevant changes you should do, trying to provide some general rules to follow.

You should monitor the SYSSTC service class velocity to get any clue of how your system is performing compared to its normal behavior from historical trends.

You should also monitor the SYSSTC service class and, where possible, the single system address spaces CPU and memory usage to identify and solve anomalies. Also in this case it would be very useful to compare the daily values with historical trends.